MUSIC LEARNING WITH MASSIVE OPEN ONLINE COURSES (MOOCS)

The Future of Learning

Learning is becoming more and more important as one of the indispensable tools to ensure future prosperity and well-being. This is the case not only for the individual, alone or as a member of a group, but also for organisational structures of all kinds. New learning paradigms and pedagogic principles, new learning environments and conditions, and new learning technologies are being tested in order to find the right combination of parameters that can optimise the outcome of the learning process in a given situation.

This book series presents to all stakeholders the latest advances in this important area, based on a sound foundation. Schools, higher education, industrial companies, public administrations and other organisational structures, including providers of learning and training services, including life-long learning, plus all the individuals involved, researchers, students, pupils, citizens, teachers, professors, instructors, politicians, decision makers etc., contribute to and benefit from this series. Pedagogic, economic, structural and organisational aspects, the latest technologies, and the influence from changing attitudes and globalisation are treated in this series, providing sound and updated information, which can be used to further improve the learning process in both formal and informal contexts.

Volume 6

Recently published in this series

Vol. 5. M. Tokoro (Ed.), Open Systems Science – From Understanding Principles to Solving Problems
Vol. 4. D. Dicheva, R. Mizoguchi and J. Greer (Eds.), Semantic Web Technologies for e-Learning

Related publications by IOS Press:

M. Tokoro and L. Steels (Eds.), The Future of Learning: Issues and Prospects
M. Tokoro and L. Steels (Eds.), A Learning Zone of One's Own: Sharing Representations and Flow in Collaborative Learning Environments
P. Kommers (Ed.), Cognitive Support for Learning: Imagining the Unknown
T. Hirashima, U. Hoppe and S. Shwu-Ching Young (Eds.), Supporting Learning Flow through Integrative Technologies
R. Mizoguchi, P. Dillenbourg and Z. Zhu (Eds.), Learning by Effective Utilization of Technologies: Facilitating Intercultural Understanding

ISSN 1572-4794 (print)
ISSN 1879-8357 (online)

Music Learning with Massive Open Online Courses (MOOCs)

Edited by

Luc Steels

ICREA (Catalan Institution for Research and Advanced Studies),
Institut de Biologia Evolutiva (UPF/CSIC), Barcelona, Spain

IOS Press

Amsterdam • Berlin • Washington, DC

ISBN 978-1-61499-592-0 (print)
ISBN 978-1-61499-593-7 (online)
Library of Congress Control Number: 2015956155
doi:10.3233/978-1-61499-593-7-i

Publisher
IOS Press BV
Nieuwe Hemweg 6B
1013 BG Amsterdam
Netherlands
fax: +31 20 687 0019
e-mail: order@iospress.nl

Distributor in the USA and Canada
IOS Press, Inc.
4502 Rachael Manor Drive
Fairfax, VA 22032
USA
fax: +1 703 323 3668
e-mail: iosbooks@iospress.com

PRINTED IN THE NETHERLANDS

Preface

Luc Steels
ICREA, Institut de Biologia Evolutiva (UPF-CSIC) Barcelona

This book is the third in a sequence of books reflecting on the future of learning, and more specifically, how computers can give support to long-distance education. [1], [2]. This challenge has recently become the focal point of attention with the rise of MOOCs that finally implement the visions that advanced thinkers proposed already decades ago.

The book has three objectives.

1. First, it introduces the phenomenon of Massive Open Online Courses, known as MOOCs. They burst on the scene of long-distance computer-based learning half a decade ago with big promises of disrupting academic education. The book discusses topics like: What are MOOCs? What is their potential? What are their historical predecessors and their future prospects?
2. Second, it presents ongoing research into making MOOCs more effective and more adapted to the needs of teachers and learners. More specifically, the book focuses on a key critical issue: Given that there are tens and even hundreds of thousands of students following a particular MOOC, how can students be given the necessary feedback during the learning process and how can their competence be assessed?
3. Third, it presents the first steps towards 'social MOOCs'. These are MOOCs that support the creation of relatively small learning communities in which interactions between learners goes beyond correcting each other's assignments. Social MOOCs try to mimic settings for humanistic learning such as a workshop, Jazz ensemble, a small choir, or a group participating in a Hackaton, in which students learn as apprentices, by solving problems, helping each other, and aided by somebody acting as tutor.

To make the discussion concrete, the book focuses on a particular domain of knowledge, namely music. Music is one of the most popular subjects (next to computing) of today's MOOCs. Many people want to learn about music, whether it is for playing an instrument, music theory, composition, song writing, or improvisation. Music requires many skills and those seriously engaged with music accept that it requires life-long learning, often taking place outside of the traditional educational system of musical academies and conservatories or private teaching.

Because of its popularity and the unique challenges that music poses to distance learning, the development of MOOCs for music has been one of the most fertile grounds for fundamental research and experimentation into MOOCs and results of experiments and the fundamental software advances that they require are already beginning to spill over into other domains of teaching through MOOCs.

What are social Moocs

Simplifying, we can say that there are two paradigms for learning and teaching: constructivist and instructional. The *constructivist approach* also known as 'natural learning' or 'humanistic learning' sees learners as active agents which autonomously explore their world by constructing rich models which they then try out while solving problems and making sense of the world. Learning within this paradigm ideally takes place within a (small) community of learners, for example, an atelier, a jazz ensemble, a fab lab, a hackathon. The teacher acts like a tutor that sets up scaffolded learning contexts and provides feedback to steer the discovery and exploration of models by the learner which are viewed as apprentices. Peers play a crucial role, both to motivate learners and to provide learning challenges and social feedback. Within the European tradition, the constructivist approach is associated with psychologists such as Vygotsky and Piaget and pedagogies worked out and put into practice by Steiner, Froebel, Montessori and, more recently, Malaguzzi. In the American tradition it is associated with Dewey and Bruner, liberal arts education, and the pedagogies of Papert or educational experiments such as Black Mountain College.

The *instructional approach* views knowledge as situation-response associations and learners as malleable acquirers of these associations through reinforcement learning. Reinforcement shapes associations in an inductive fashion through positive and negative examples and through reward and punishment administered by an external agent (possibly the environment). The instructional approach is associated in the US with psychologists such as Thorndike or Skinner and with pedagogies based on strict lesson plans with continuous assessment and clear reward and punishment. The instructional approach is also the foundation for connectionist inductive learning systems developed in AI.

The instructional approach has been shown to be very effective for the acquisition of basic skills and allows clear assessment and standardised education. It is therefore often imposed on teachers through central educational bureaucracies. But it is also known to lead to severe problems such as demotivation of both teachers and learners.

Both approaches have been used in computer-based education. For example, the programmed instruction method invented by Skinner can easily be turned into a computer-based instructional system, and some MOOCs, including many of the current MOOCs for music, follow this model rather closely. It was less obvious at first to use computers for supporting the pedagogy of open-ended constructivist learning, particularly in the context of distance education. However there have been some early significant developments showing the way. The best example is Seymour Papert's LOGO programming environment, in which learners discover mathematical concepts through programming the movements of a Turtle.

Based on this experience and the growing penetration of computers and smart phones in our daily lives, several papers in the present book argue that it is precisely through the use of computers, in particular through social media and computational means for assessment and tutoring, that a constructivist pedagogy can be put into practice on a much larger scale, leading to the promise of *social MOOCs*: MOOCs that try to foster a community of learners who share their work and help each other through feedback and cooperative problem solving.

Although traditional forms of classical music training, particularly the mastering of an instrument, often follow a Skinnerian instructional approach (piano teachers hitting their students on the hands is not uncommon), there is a consensus that greater enjoyment, motivation, and more powerful learning takes place within a constructivist peer setting, for example a string quartet or a small Jazz orchestra where players scaffold and motivate each other, progressively enhancing their skills. Social MOOCs try to recreate such learning conditions and thus make natural learning available to students who do not have access to the kind of peer and apprentice-style tutor interaction assumed by humanistic learning.

Structure of the book

The papers in this book all discuss steps towards social MOOCs: their foundational pedagogy, platforms to create learning communities, methods for assessment and social feedback, and concrete experiments. These papers can be read on their own, but they also strongly relate to each other and are presented in a logical progression, from background and pedagogical theory to concrete platforms and experiments. The papers are organized into five sections: I. Background, II. The role of feedback, III. Platforms for learning communities, IV. Experiences with social moocs, and V. Looking backward and looking forward.

Part I. Background

The first part of the book provides the basic background to later papers. It starts with a paper by LUC STEELS: *The coming of (social) MOOCs*. He introduces the notion of Massively Online Open Courses, sketches how MOOCs arose to deal with the 'crisis in education', how they rapidly spread thanks to the Internet, and what the current state of deployment is. The paper also introduces the concept of social MOOCs, why it is important to develop them, and which obstacles need to be overcome.

The second paper by JOHAN LOECKX entitled *Learning music online*, surveys the state of the art in on-line music learning. It provides background for the case studies reported later in the book and is intended as a guide for teachers and platform designers. Loeckx not only surveys and classifies the intense ongoing activity in online music learning but also examines critically some of the issues with available systems and argues for a better grounding of online music learning in pedagogy.

Part II. The role of feedback

Part II lays out the pedagogical foundations for social MOOCs, focusing particularly on the issue of feedback, not in terms of rewards and punishment but feedback as required to create stimulated open-ended learning environments that support constructivist learning.

The first paper by MARK D'INVERNO and ARTHUR STILLS entitled *Social feedback as a creative process*, sketches the historical roots of constructivist and instructional pedagogies within the Anglo-American tradition, and introduces the notion of social feedback as a key ingredient for social MOOCs.

The second paper by LUC STEELS, entitled *Social Flow in Social MOOCs*, pulls the concept of flow out of its traditional individualistic character to examine its potential role in the creation and sustainance of a motivated learning community within the setting of MOOCs.

Then there are two papers which give very personal accounts of learning trajectories for music. They both illustrate the importance of humanistic learning for the development of top musicians. The paper by RAY D'INVERNO, a renowned Jazz pianist, is entitled

Teaching Jazz improvisation: a personal experience. It sketches his learning trajectory, recounting how the setting of a small Jazz ensemble which interacted regularly with experienced players, played a key role for him becoming a Jazz master and how he has tried to translate these insights into a teaching methodology recently used as the basis of a MOOC.

The paper by JOSEP-RAMON OLIVE, an upcoming opera singer. It is entitled *Learning to be a singer* and sketches his personal learning trajectory. He emphasizes again the importance of a humanistic education, where peer activity and guidance by a tutor create learning opportunities without a fixed curriculum or a rigid reinforcement framework.

Part III. Platforms for learning communities

Current MOOCs act mostly as content delivery platforms to be used by an individual learner with no direct social contact with others. The main interaction happens anonymously when learners are asked to correct some of the assignments of other students. Many MOOCs do feature some social media facilities (such as forums) and encourage physical encounters between other students in the same area. But the work on social MOOCs discussed in this book go much further, proposing and experimenting with platforms for the creation of learning communities where individuals are no longer anonymous but interact intensely with each other.

The first paper, entitled *Music circle: Designing educational social machines for effective feedback*, by MATTHEW YEE-KING, MARIA KRIVENSKI, HARRY BENTON, ANDREU GRIMALT-REYNES, and MARK D'INVERNO introduces a social MOOC platform for music learning that incorporates both the technologies for sharing music and mechanisms for supporting social feedback. The paper describes the participatory design methodology used to conceive this platform and reports the results of extensive evaluation studies with real users.

The second paper, entitled *Giant Steps in Jazz Practice with the Social Virtual Band* by MATHIEU RAMONA, FRANCOIS PACHET and STANISLAW GORLOW, describes another example of a social MOOC platform. It is geared to learn about Jazz improvisation, recreating the kind of interactions one sees in a small Jazz ensemble. The system not only integrates facilities to play along with standards, to store and share the results of these practice sessions, and to create or accept feedback with peers, but also tools for automatic machine-based feedback, for example for playing scales.

The third paper is entitled *Steps towards intelligent MOOCs* and contributed by KATRIEN BEULS and JOHAN LOECKX. It explores another aspect of humanistic learning, namely tutoring, using musical composition, specifically the writing of counterpoint, as a case study. The paper argues that the methods and technologies developed in intelligent tutoring systems can be integrated in MOOCs so that students get much more sophisticated feedback. The paper also reports on experiments to put flow theory at the service of scaffolding challenges for students.

The final paper of Part III is entitled *Collaborative Peer Assessment using PeerLearn* by ISMEL BRITO, PATRICIA GUTIERREZ, KATINA HAZELDEN, DAVE DE JONGE, LISSETE LEMUS, NARDINE OSMAN, BRUNO ROSELL, CARLES SIERRA and CARME ROIG. It proposes a platform which has the creation and management of lesson plans devised and overseen by teachers as its core and then adds facilities for peer assessment. This platform is intended to support blended learning, in which distance education is integrated with traditional student-teacher interaction.

Part IV. Experiences with social MOOCs

This part of the book documents concrete efforts to use social MOOCs with 'real' users, both in the context of organized education (blended learning) and in the context of an open audience solicited through the web. The first paper entitled *Using social media to revive a lost apprenticeship model in jazz education* has been contributed by ED JONES, a renowned Jazz saxophonist, and HARRY BRENTON. They used the Music Circle Platform (introduced in Part III) to support a course on Jazz saxophone, studying in particular the role of tutor and peer feedback, whether a MOOC-like environment could help to revive the apprenticeship model traditionally used in Jazz education,

The second paper is entitled *Improving music composition through peer feedback: experiment and preliminary results*, contributed by DANIEL MARTIN, BENJAMIN FRANTZ and FRANCOIS PACHET. It reports on an experiment using the Virtual Social Band MOOC environment, raising two questions: (i) To what extent can peer feedback affect the quality of a music composition? and (ii) How does musical experience influence the quality of a feedback during the song composition process?

This part of the book ends with an intermezzo: a personal account by FIAMMETA GHEDINI on following a MOOC in the form of a comic strip entitled 'So, I've been following a MOOOOC'. The MOOC was about song writing and offered by the Berklee College of Music (Boston).

Part V. Looking backward and looking forward

The final part of the book puts MOOCs in a broader context. KEN KAHN, one of the early pioneers of the creative use of computers in education, discusses in his paper *A half century perspective on the role of computers in learning and teaching* the development of learning environments, and particularly open-ended constructivist learning environment such as LOGO Mindstorms devised by Seymour Papert.

GEORGE VAN DE PERRE, who is one of the early pioneers in distance-education and on-line learning, discusses in his paper *Blended learning and MOOCs* basic issues for the introduction of MOOCs within the context of traditional universities, and sketches the current movement towards blended learning that exploits the novel opportunities of MOOCs but integrated within the existing university framework.

CONCLUSIONS

The following key conclusions can be drawn from the papers in this volume: (i) MOOCs have arisen as a logical consequence of marrying long-distance education with the web and social media. They are here to stay and provide a valuable addition to the toolkit of learners and teachers alike. (ii) Most MOOCs today are based on instructional pedagogies and content delivery through the web, but there is the opportunity to build a new generation, which we call social MOOCs, that supports more powerful humanistic learning, which is much more adapted for many domains, such as Jazz improvisation. (iii) Assessment and feedback play a crucial role in all pedagogies and it is critical also in the development of social MOOCs. We need novel approaches, such as peer feedback, in which learners assess each others' achievements, automated assessment, in which algorithms take over some of the basic checks in a student's work, intelligent tutoring, which not only identifies errors but also suggests ways to repair them, etc. Several papers in this book show very concrete examples on how this can be done and report experiments testing whether these proposals work out in practice.

Due to the rapid advances of knowledge, globalization, budget cuts, and a bombardment of information, it is not easy for current generations to still find the opportunities, time, and focus for profound learning, even though this is more than ever necessary to survive in today's stressful economic climate. Technology is not a panacea for fixing the enormous challenges facing today's educators and learners, however we hope that the technologies developed here can lead to more effective and more humane learning opportunities for a larger group of students.

Acknowledgement

This book and most of the research reported here came out of the European FP7 project PRAISE (EU FP7 number 388770), funded by the European Commission under program FP7-ICT-2011-8. The seeds for many of the papers published here are the outcome of a workshop organized by Luc Steels with the assistance of Emìlia Garcia Casademont in Casteldefells (Spain) on 10-12 December 2012. I am indebted to Jorge Piz Dico (UPF, Barcelona) for proofreading most articles and helping out with Latex type setting issues, and Maria Ferrer Bonet (UPF, Barcelona) for editorial help and the handling of administrative matters. I also thank the people at IOS Press, particularly Carry Koolbergen and Paul Weij, for their efficient handling of this publication.

References

[1] Tokoro, M. and L. Steels (eds.) (2003) The future of learning. Issues and Prospects. IOS Press, Amsterdam.

[2] Tokoro, M. and L. Steels (eds.) (2004) A learning zone of one's own. Sharing representations and flow in collaborative learning environments. IOS Press, Amsterdam.

Contents

Part IV. Experiences with Social MOOCs

Part V. Looking Backward and Looking Forward

Part I

Background

This page intentionally left blank

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-3

Chapter 1. The Coming Of MOOCs

Luc Steels
ICREA, Institut de Biologia Evolutiva, Universitat Pompeu Fabra and CSIC

Abstract. A MOOC is a Massively Online Open Course. It is massive because there are many students (sometimes hundreds of thousands). It is online because it uses the Internet for course delivery. It is open because it is publically available to anyone without selection barriers or payment. And it is a course, teaching a particular subject, often in engineering and science, but increasingly in all domains of human knowledge including the arts. MOOCs burst on the scene of online distance learning in the fall of 2011 and caused a wave of excitement followed rather quickly by a wave of scepticism and resistance. What are MOOCs? Will they help to deal with the ‘crisis in education’? How do they fit within the earlier developments in distance-education and the use of computers and telecommunication for supporting learning processes? What are the limitations of MOOCs? How can we strengthen them and fully profit from their potential? This paper addresses these questions from my personal viewpoint as an educator involved for decades in teaching and online distance-education. It looks at MOOCs, bringing in a European perspective, and suggests avenues for further research and practice.

Keywords. MOOCs, online learning, distance-education, social MOOCs.

Introduction

For several decades I have been a university professor at the Free University of Brussels (VUB) and thus active in educational practice, at first teaching ‘ex cathedra’ for large groups of students in the first bachelor years of computer science, and gradually focusing on master’s level courses in Artificial Intelligence and directing PhD’s with more than 30 having graduated so far. In addition, I have been active throughout my career with distance education.

During the late eighties and early nineties I worked with the Dutch Open University to turn some of my courses into educational materials that could be studied at home with occasional tutoring in the educational centers of the (Dutch) Open University. The most successful course on knowledge engineering [23] was attended by thousands of students, which is a substantial number given that the course was in Dutch and so the possible population of students was restricted to the Netherlands and Northern Belgium. I was also involved in a start-up company called Knowledge Technologies, located in Brussels, that specialised in the application of knowledge-based systems to education. The company later became part of Didael, still an italian leader in distance education (`http://www.didaelkts.it/`).

In the mid-nineteen nineties, I started to work extensively with EuroPACE, a European-wide organisation for distance education [28]. EuroPACE was using primarily televised lectures, broadcast via satelite or through national broadcast organisations. The World Wide Web started with a few servers only in 1992 and so satellite television was the only option at that time to spread video materials. For EuroPACE, I made a total of 60 hours of lectures about a variety of topics, many in the domain of computer science, more specifically Artificial Intelligence.

One of my biggest projects, in 1994-1996, was titled *Science on the Edge of Chaos* (`https://arti.vub.ac.be/previous_events/chaos/intro.html`). It was anounced as “an interactive multi-media web on complexity and chaos”, which was at that time an emerging hot topic. Different themes were covered through the eyes of complex systems science - from physics, biology, and technology to cognitive science and socio-economics. Each theme was worked through five components:

1. A lecture introducing basic concepts from complex systems science in small chunks, using a format similar to today’s MOOCs.
2. A case study based on work of a leading scientist. For example, one case study focused on chaos in the solar system based on the work of Jacques Laskar (Observatoire de Paris)
3. A conversation with a high profile scientist on chaos and complex systems science (including Noble prize winners Ilya Prigogine, Manfred Eigen and Christian De Duve).
4. A live panel discussing hot topics and recent progress for each theme.
5. Various web resources and ways for student interaction using email, videophone and occasionally live telephone interaction during the panels.

This project was, in 1994, probably the first one in Europe to use the web as a rich educational medium, and it clearly foreshadowed many aspects of the MOOCs that arose twenty years later. No video could be transmitted through the web and so television broadcasting had to be used. The number of people having access to the web at that time was still very limited and also financial resources for the independent development of distance-education were hard to come by, so this particular project stopped after six themes were covered. However, it illustrates clearly that the vision of online education using video lectures and web resources was already alive in the nineties and it was just a matter of time—until personal computers, widespread internet access, and sufficient server capacity were available—before it would become really possible and scale up to the MOOCs we see today.

My various educational activities and experiences with online distance education in the eighties and nineties culminated in a series of workshops in the beginning of the 21st century, which I organised in Bagnols (France, 2002), Calheiros (Portugal, 2004) and Carmona (Spain, 2003). The outcomes of these workshops were published in two books on *The Future of Learning* [25], [24]. The workshops were motivated by two insights: a growing sense of a crisis in education, and hence pessimism, but also a growing awareness of the opportunities that new technologies, particularly those centered around the use of computers and telecommunications, were beginning to offer, and hence optimism. The current rise of MOOCs is realising many of the visionary ideas that we were dreaming about decades earlier.

The rest of this paper first discusses the nature of the crisis felt in education. I then turn to the trend towards personalisation and decentralisation that is emerging as a possible answer to deal with the crisis, and discuss how the development of information and communication technologies is progressively enabling this trend to become reality. Against this backdrop, I then focus on MOOCs and reflect on their positive sides but also on the criticisms beyond the initial hype.

1. The crisis in education

The worrisome problems encountered by the educational systems in European countries were already jumping from the facts around the turn of the century, and they have been since exacerbated.[1] What is the nature of this crisis and what are the possible causes?

1.1. The nature of the crisis

Here are some of the key observations [6]:

- There has been a clear *diminishment in academic achievements*. In France, 21 to 42 % of the students finish their high school education without adequate reading and math skills, which is lower than in the 1920's. Still in France, there is an increase in the proportion of low achievers from 15.2 % in 2000 to 21.8 % in 2006, showing that the downward trend is continuing [16]. Similar statistics can be found for other European countries. The effect of diminishing achievement is percolating into the higher educational system. For example, many university students, even in scientific and technical fields, still graduate without adequate programming skills or the ability to express themselves in clear writing. They have memorised some of the main facts in their fields but are unable to engage in the kind of life-long learning that is required in today's rapidly evolving society.
- There has been a growing trend towards *violence* in schools. For example, in France 15,000 incidents are reported every month (already in 2002). Some of these incidents—occasionally caused by parents supposedly 'defending' their children—have lead to the death of teachers or significant physical harm [21]. This trend is obviously making it difficult to establish the peaceful surroundings needed for real learning and is strongly demotivating excellent people to choose the educational profession.
- The number of *student dropouts* has been increasing steadily. They stand at 30 % in Italy or the Netherlands for university education, although they are lower in countries with stricter entrance requirements, such as the UK where it is 16 % [18]. This high rate is taking away resources needed for the students that strive for excellence and hence it is lowering the quality of education.
- There is growing *teacher unrest*, with national strikes and many teachers leaving the system early in their careers. There are growing difficulties to recruit new teachers, particularly for mathematics, computing and science, with significant shortages, even in countries such as Germany where the educational system is under less stress [4]. One side-effect of this trend is that courses (for example

[1] Here are two interesting data sources for education in Europe and comparisons to the rest of the world: http://ec.europa.eu/eurostat/web/education-and-training/overview and http://www.nationmaster.com/country-info/stats/Education

computing) are taught by teachers who have hardly any knowledge nor training in the fields that they are teaching.
- There are persistent complaints from the industry that they cannot find enough people with the right qualifications. Particularly in the domain of science and engineering the situation is dramatic. For example, in Germany alone there are 14 000 vacancies for electrical engineers. Not enough qualified students are graduating and those that graduate are lacking essential skills.

Despite these undeniable trends, and despite significant expenditure (for example, on average 7780 EUR is spent per student per year in France), many educational reforms have failed. Some of them because they were ill-conceived, but also because there is often resistence from teachers, who are increasingly getting tired of bureaucrats with little knowledge about what it means to stand before a classroom meddling with educational practice and causing unnecessary disruptions by the imposition of harmful policies.

The sense of a crisis is very much with us today, and for many observers the situation is only becoming worse. Whereas initially only inner city schools in the big cities were affected, problems have now spread to smaller towns and to all European countries. Moreover opportunities for ‘after school’ education in music, arts, sports, languages, etc., which were an essential part of the European public educational system in the past, have began to dry up due to budget cuts, thus giving fewer chances to students to complete their education with skills and knowledge that used to complement the standard curriculum, which is increasingly oriented towards vocational training only.

1.2. Possible causes for the crisis

The causes of the crisis in education have been much debated. The educational system is certainly partly under heavy stress due to the significant increase in the number of students and the evolution towards a multi-cultural society in Europe. But there are some other culprits that are often cited:

1. The European educational tradition has always been based on the laudable ideal that anyone should be given maximum opportunity for reaching the highest possible level of self-development and that this should not be based on the background or financial resources of the parents. This ideal plus significant population growth have lead to a massive increase of students, particularly in higher education. For example, the bachelor degree in France has jumped from 4,9 % of the total student population in 1950, to 19 % in 1970, 25 % in 1980, 43 % in 1990 and 62,7 % in 1995. This massive increase has had two side effects: (i) It is hard to believe that all these students have the skills, motivation and talents required, so that unavoidably the dropout rate becomes higher. (ii) The scarce teacher resources are spent on trying to educate at all cost a large group of students who are not at the appropriate skill level and have no interest in learning, taking away resources needed for reaching excellence with the other students. Larger classes have lead to a degradation of teacher-learner contact, causing problems of discipline and student demotivation.

2. The trend towards “*everybody has to have a university degree*” has been reinforced by the European Commission’s mantra of the ‘knowledge society’, claiming that everybody has to become a knowledge worker[26]. The overemphasis on ‘academic knowledge’ has lead to a scarcity of skilled plumbers, carpenters or manufacturing workers but an overabundance of people with ‘higher’ degrees—often of doubtful value—but

without any prospect of work. It has accelerated the demise of manufacturing in Europe, a high unemployment rate, and an increasing number of young people working below the level supposedly granted according to their degree. It has also reinforced the sense of crisis in education, where schooling no longer prepares you for later professional and cultural life.

3. The Bologna reorganisation of the higher educational system has had a very negative impact on the level and quality of teaching in several European countries. The reorganisation was again based on a laudable objective: Leveling the different paths of higher education (e.g. between technical engineering schools and academic engineering faculties) and thus giving greater opportunity to all. In practice it meant that students with educational programs with lower challenge levels were upgraded (a lot of vocational engineering and arts degrees were from one day to the next declared to be at a university level), but, unavoidably, the programs that targeted higher challenge levels, specifically academic degrees, were downgraded because students were allowed to move between degree programs practically without entry barriers.

4. The Bologna reform also made students more responsible for their own educational trajectories and gave them enormous flexibility to choose. This has lead to student classes with very unequal backgrounds and competences, making it difficult to keep the necessary level of excellence. In the end, it has become almost impossible to ensure that students have the level of competence supposedly guaranteed by their diploma. Many students now slide through the system, seeking out the 'easiest' way to reach the desired diploma and avoiding the hard work that would normally be required. Attempts to bureaucratise this process (specifically through the European Credit Transfer and Accumulation System [http://ec.europa.eu/education/tools/ects_en.htm] set up in the Erasmus framework) are not satisfactory because they assume that it is easy to define learning outcomes and that education can be easily decontextualised.

5. The central role of the state in education started in the 19th century in Europe. Education became compulsary with state funding provided directly to state schools or indirectly to religious organisations that set up 'private schools', which were however to a large extend still controlled by the state [5]. Curricula became standardised and levels and requirements were nationally defined. Teacher education and qualification became centrally organised and regular quality controls were introduced in the schools. This has certainly lead to a tremendous improvement of education for the largest possible segment of the student population.

However, the centralised state-wide management of education has also lead to a very slow adaptation to the profound changes in the educational needs of society (for example a very slow inclusion of computing skills in standard curricula). There is a growing mismatch between the needs of a particular community or of the industry in a particular region and the national objectives. Clearly, a highly multi-cultural and multi-lingual class in an urban environment should be allowed to function in a very different way from a class in a rural mono-cultural context. Technical schools in a region with a lot of industry specialised in mechanical engineering need to cater to different learning objectives compared to a technical school in a region specialised in the financial service sector. The mismatch between local needs and global state-imposed requirements is often cited as one of the causes for a high drop-out of students who feel curricula and pedagogical methods are not adopted to their needs.

2. Personalisation and decentralisation

There is certainly no magic wand that will make all these problems go away, and top-down bureaucratic intervention has only seemed to make the situation worse. Every European country has tried to combat the crisis in education in different ways and inspirational educators have taken many, often very valuable, initiatives. Generally speaking, we see a trend towards *personalisation* and *decentralisation*.

Personalisation means that independent actors (private companies or governmental institutions) are stimulated to provide educational opportunities and the students themselves (with their parents) are stimulated to actively choose and shape their educational trajectories for their own good and adapted to their own needs. As we will see, MOOCs fit entirely within this trend.

Decentralisation means that decision-making is no longer entirely controlled by a central educational bureaucracy but devolved to local governmental organisations who can operate within an envelope of financial means and requirements provided by the state [32] or even to small groups of individuals that band together to learn about a particular topic. Local actors are in this case regions, cities, networks of universities and schools, or semi-governmental non-profit organisations.

2.1. Provisioning by private actors

One radical response to deal with the crisis has been to introduce in Europe a market approach to education, thus following the American example [29]. This approach is being tried out most strongly in Britain, which, since the Thatcher era, has veered towards an American style capitalist system that relinquishes the role of the state as much as possible in favor of the private sector.

The basic principle of this model is to stimulate private actors to produce educational materials. Students and their parents are viewed as consumers that buy education the same way they buy a car or a smart phone. The providers go into competition with each other, both to attract students and to prove the worth of their diplomas with employers, which in turn ensures (at least in theory) that future students select the best provider. A school that produces successful graduates will attract more or better students and thus get the income to sustain itself. According to the market logic, competition should lead to an improvement of what is on offer, greater efficiency in creating or supplying education, and an improvement of consumer choices, i.e. it should force students to choose more carefully what they want to study based on their talents and the prospects of a later job.

Today there is certainly a willingness to consider contributions from the private sector in education. In fact in some domains private tutoring has always been very common. For example, music teaching is traditionally done by private teachers that interact with students on a one-on-one basis and if music education is organised by the state, it follows closely a highly personalised pedagogy. Intensive sports training is also done within the context of private clubs and with trainers that give a highly personalised attention to talented individuals. Many companies organise their own continuing education to keep their workforce up-to-date and highly skilled and most of them make heavy usage of external companies that are specialised in offering specialised education.

In all European countries we see a growing trend towards public-private partnerships, for example to provide and take care of schoolbuildings. There is also a timid rise in fees for higher education, except in the UK where the rise in tution fees has been dramatic. There is also a growing number of private actors which are providing some form of higher education, as the ‘monopoly’ of state-funded institutions to hand out certified degrees has been broken in several countries, such as Belgium for example. So these developments are paving the way for a greater involvement of private actors, such as MOOC providers, in higher education.

But the application of the market principle to education on a grand scale is seen in most European countries as a bridge too far. Here are some of the problems:

- As in any capitalistic market-based system, the value of a product is translated into how much it costs and private actors only invest if there is a profit to be made, i.e. if the price of the product is sufficiently high to cover at least the cost. This poses a problem for education because there is a basic unavoidable cost for the knowledge-intensive teaching model that is generally believed to be necessary for true learning and knowledge creation. So the price for education, assuming no state support, is necessarily high.
  Such a high cost means in practice that the financial means of the parents determine the quality of the education of their children or that students have to take on loans which they have to pay back during a significant part of their life, trapping some of them into a cycle of poverty. Privatisation of education in the UK is already beginning to show signs of this. The average price of the university tuition fee in Britain is now 12,000 eu/year. Although this may still sound reasonable compared to American university tuition, which goes up to 50 000 EUR/year, it is still viewed as extraordinary high from the viewpoint of Europe, where parents pay on average between 500 and 1 000 EUR/year and the state funds the rest. Moreover, even if the US private-sector model has lead to elite universities, “*among developed countries, the United States is 55th in quality rankings of elementary math and science education, 20th in high school completion rate and 27th in the fraction of college students receiving undergraduate degrees in science or engineering.*” [12]. In other words, privatised education reinforces existing inequalities and thus not ensure a sufficiently high educated workforce.
- It has proven difficult to create a ‘market for education’ in Europe, partly because the traditional providers (i.e. universities and technical schools) all ask the same price when they are given the right to set the price themselves. Moreover the choice by the student is often not based on supposed quality (which in many countries is guaranteed by the state to be the same across all institutions anyway) but on other factors, such as proximity or family tradition.
- Universities have a broader task than education. They contribute to the common good through research and service to society and this is generally acknowledged to underpin the economic viability and societal functioning of the regions in which they operate. By pushing universities to compete in an optimised and privatised educational market, many critics believe that we will see a ‘tragedy of the commons of knowledge’, i.e. a general shrinking of the knowledge base from which companies and government institutions can draw.

2.2. Provisioning by governmental actors

The trend towards personalisation is not only motivated by adepts of a liberal and capitalistic organisation of society. In the seventies there were educational theorists, such as Ivan Illich with his influential book “Deschooling society” [10], which were advocating the same, except that they believed that provisioning should happen through public channels rather than by privately owned for-profit companies. Illich argued for example for the need to create computer-based learning webs, where individuals could access reference services that describe educational opportunities for formal learning, and skill exchanges, which are peer-to-peer services where those with skills offer to teach others. Many of these ideas are beginning to become reality in the context of MOOCs.

Governmental and non-profit infrastructures that come closest to providing personalised educational approaches are the ‘open’ universities that were created in the seventies. Open universities allow students to take single courses independently and at their own pace. They thus provide access to higher education for those who have to (or want to) mix education with work. It creates new opportunities for continuing education after a degree, and allows those who dropped out from a traditional courseprogram to restart their educational trajectory.

The largest example of this kind is the Open University in the UK, but there are similar institutions everywhere in Europe, such as the Fernuniversität Hagen in Germany, the Open Universiteit in the Netherlands, the Universitat Oberta in Catalunya, Uninetto in Italy, etc. The ‘open university’ courses are often produced in collaboration with professors and their teaching assistants at existing universities. Additional resources are provided by a central organisation that helps to define the course objectives and assists in designing and developing tutoring materials, exercises, and evaluation methods that can be used without physical proximity.

The open universities offer certified degree programs which are equivalent to those offered by ‘brick-and-mortar’ institutions. This implies that they have strict assessment procedures comparable to regular universities. They include personal tutors that provide also academic expertise, guidance, and feedback, although students are primarily supposed to work on their own and at their own pace. There are opportunities for social interaction with others, possibly through online media, study networks and course forums, as well as through occasional physical meetings at open university centers. Students pay per course. Although the price is usually lower than that of the normal universities, it is still much higher compared to the almost free education on offer in MOOCs.

Open universities are clearly the real pioneers of distance education and millions of students all over Europe have participated in their courses with great satisfaction. They have worked out novel pedagogies and techniques for courseware that are (or should be) of great value to MOOC developers. Open universities have also pioneered many of the information and communication technologies for distance education on which MOOCs are based, as discussed more fully in the next section. Open universities have proven most valuable for continuing education and they have certainly not replaced the traditional higher education institutions that are still teaching the bulk of students. Moreover the open universities draw for a lot of their course material on ‘brick-and-mortar’ universities.

3. The role of Information and Communication technologies

Information and communication technologies (ICT) have progressively come on the educational scene since the mid-20th century. It started with the ‘school radio’ in the fifties and the 'school television' in the sixties. Initially most households did not have radio or television and students listened to the radio or watched broadcasts inside classrooms, providing examples of so called *blended education*, but as radio and TV became consumer products, and audio and video playback became widespread, educational technologies for self-teaching started to take off and were adopted enthusiastically by the open universities or other organisations such as EuroPACE offering continuing education.

The next step came with the introduction of computing technologies, gradually slowly becoming available in classrooms and school administrations in the seventies and then in the home in the nineties. At the same time, the Internet took off, again at first slowly in the seventies, but then progressively propagating like a huge wave, shortly after the World Wide Web protocol made information transfer and access straightforward and scalable.

Today ICT has in many ways become crucial in education. First, the most basic impact comes from *Learning Management Systems*. They are used by schools and universities to organise the curricula, deliver content, identify and assess learning goals, and collect data about the individual or a class. They are also used by students themselves in order to register for courses, download course materials, upload exercises, and see their grades. MOOCs had to support all these activities, but they needed to scale up the technology of Learning Management Systems to handle hundreds of thousands of students.

Second, attempts have been made from the sixties onwards to go much beyond administrative purposes, and to build systems that assist in the educational process itself. A whole series of information and communication technologies have resulted from these various efforts. Most of them have often not gone beyond small-scale academic experiments, although there are quite a few companies that are marketing already concrete applications.

The developments towards the use of computers for teaching can be grouped into three main trends: Computer-assisted instruction, intelligent tutoring, and open learning environments.

3.1. Computer-assisted Instruction

The first use of computers for teaching, already in the sixties, applied the *programmed instruction* pedagogy developed by behaviorist B.F. Skinner. This pedagogy strictly streamlines course materials in small chunks and suggests students to incrementally go through units, each time doing the exercises associated with a small chunk. Students have to check their own answers and can only advance further after answering correctly. Soon richer forms of computer-assisted instruction were developed, such as the PLATO system, where PLATO stands for Programmed Logic for Automated Teaching Operations. In addition to computer-assisted instruction in the Skinnerian style, PLATO grouped functions that we now find in Learning Management Systems, such as managing course content and student participation. The system also allowed the inclusion of a variety of additional learning materials, such as texts, videos, computer simulations, and it provided 'courseware' tools for creating online lessons [22]. The PLATO system is therefore an important forerunner of the MOOC.

Many of the currently successful computer-based teaching systems, such as Rosetta stone (`www.rosettastone.com`) for language learning, are still following the pedagogical format pioneered by Skinner's learning machines. And many large-audience MOOCs, such as the ones for Artificial Intelligence and Machine Learning developed by Udacity and Coursera, are essentially based on the same behaviorist pedagogical principles offering as well the added functionality of early pioneers like PLATO, but now made more attractive with better interfaces, faster multi-media savvy computers, and delivery through the internet rather than through floppy disks or CD-ROMS.

3.2. (Intelligent) Tutoring Systems

When Artificial Intelligence started to mature in the seventies, the objectives of developers became more ambitious [17]. They started to target systems that had three capacities: (i) to be (automated) problem solvers themselves in the domain in which they were tutoring, for example, be able to parse or produce sentences, solve mathematics problems, compose music, (ii) to build a student model that would explain student errors, and (iii) to actively plan a course of action based on particular tutoring strategies. Although a lot of work was done by many AI researchers, the technology needed to achieve intelligent tutoring was—and still is—too complicated for widespread adoption. But it is clear that if online education through MOOCs is ever going to come anywhere close to human teaching, it will have to integrate aspects of intelligent tutoring [1].

3.3. Open learning environments

The Skinnerian behaviorist pedagogy is most effective for a particular type of knowledge, such as the rote learning of words. But many psychologists, such as Piaget and Vygotsky, and educators, from Steiner to Malaguzzi, have argued that for many subjects learning is much more effective in an open environment in which learners are offered a variety of materials which they can explore and from which they can discover solutions themselves. In the mid-seventies, several researchers started to investigate how computers could help to put this constructivist learning approach into practice. One of the main advocates was Seymour Papert, who devised the LOGO programming language and Turtle geometry as open learning environments for mathematics and computing [15].

The most well-known MOOCs so far tend to lean towards programmed instruction and therefore stick to rather rigid educational lesson plans and immediate testing, but it is possible to conceive of MOOCs that pursue a constructivist pedagogy, for example, by providing rich simulation environments that stimulate students to construct in a game-like environment their own models of particular phenomena and test them out [11], or by heavily integrating social media to make up for the isolation of the individual learner and thus stimulate collective knowledge construction [3], [14].

In fact, the term MOOC was originally used to describe experiments in 2008 by George Siemens, Stephen Downes and others in Canada to implement a so called ‘connectivist’ pedagogy [20], that is closer to the constructivist learning in open environments that Papert had been advocating than to the programmed instruction approach, although it targets other subject domains, primarily in the humanities, such as thinking or writing.

From the viewpoint of the connectivist approach, content within some domain of knowledge (i.e. facts) is not the most important thing that needs to be learned, because this content is shifting rapidly and abundantly available today through web resources such as Wikipedia. Rather learners should acquire the skill to access, aggregate and navigate through content and teachers are facilitators that should guide this process.

The MOOCs built with a connectivist foundation are now called connectivist MOOCs or cMOOCs to distinguish them from the behaviorist MOOCs (now sometimes called xMOOCs) that follow primarily the programmed instruction format.

Instead of tutors, the connectivist pedagogy talks about facilitators. Students are presented with an overwhelming amount of materials (instructions, texts, interviews, videos, blogs, etc.) through which they have to browse and skim through, just like you have to do when exploring a new subject through the web. They are expected to bring order and aggregate these materials, remix them, and feed them back to others through blogs, email, streaming, chat channels, message boards, screen sharing, and blogs, as well as Internet-driven social media such as Twitter and Facebook. In the next step students are expected to go beyond the knowledge already out there and present and share it, thus adding to the big pool of content accessible to the other students so that the cycle can start again and new connections and content can be learned and created. Advocates of the connectivist pedagogy argue that this mode of operation teaches students to participate in life-long learning communities which are essential in this fast moving digital age.

4. The rise of the MOOC

Although we have seen that there have been many precursors for the development of MOOCs, they grabbed world-wide attention only in 2011. What sparked world-wide attention was that MOOCs not only addressed a key challenge which was not explicitly addressed before by online learning systems, namely how can online courses be designed and distributed so as to reach a very large number of students using a very small team of educators (typically one professor with technical assistance), but that they actually achieved their goal of attracting and coping with these high numbers of students (usually tens of thousands and sometimes a lot more). The challenge to deal with such a high number has required several important novel technical and pedagogical ideas on how education can be structured and delivered online.

The key properties of MOOCs are as follows [8]:

- Most MOOCs consist of a series of very short video lectures (much shorter than a normal class) immediately followed by exercises that can be checked before moving on, thus going back to the Programmed Instruction pedagogy. There are additional supporting materials (for example articles, software, links to web resources) and final assignments that are graded.
- Grading is done by other course participants or by computer programs in order to cope with the huge number of participants.
- MOOCs are in principle delivered free of charge over the internet and there are no admission criteria. Students who sit out the complete course and do all the assignments receive a certificate of participation, often after paying a fee.
- MOOCs are participatory: They have facilities for bringing students together so they also create a social context, which has been shown to be important for motivating learning. Like social media, MOOCs try to establish links between individuals, creating foundations for life-long future interactions, similar to the way that class mates in residential schools form a social network that is important for the rest of their life.

- MOOCs collect student data that is then used by data mining algorithms to study and improve learning efficiency or student engagement [2].
- Although a MOOC is said to be 'open', this does not mean open in the sense of open software, but rather that all materials and student discussions are open for everybody. For the rest, MOOCs are entirely closed. They are designed by a professor and his or her team and they cannot be altered or distributed by anyone else. MOOCs provided through the Free Technology Academy (http://www.ftacademy.org/about) are an exception. Its courses are open, in the same sense as open software: They can be edited and remixed by anyone for further distribution.

MOOCs are not based on any top-down state initiative. They introduce innovation entirely in a bottom-up disruptive fashion. Most MOOC providers subscribe to the competitive market model of education. Because courses from suppliers are publically available, they can be compared and presumably the reputation of the best courses will progressively attract more students, the same way better products become dominant in the market place. Although courses are now mostly free, this might change as the venture capitalists that currently fund MOOCs request the (high) returns on their investment they are accustomed to.

4.1. The starting phase

The first MOOC that attracted a very large audience appeared in the fall of 2011. It was based on the Introductory Artificial Intelligence course taught at Stanford University by Sebastian Thrun and Peter Norvig. The course setup was quite similar to courses offered by the various open universities or televised course providers such as Uninetto in Italy or EuroPACE. What was new was the sudden influx of an unexpectedly huge number of students: 160,000 took the course, purely based on hear-say through the web.

This success was partly due to the fact that the course was already well-known beforehand, as it was based on an earlier Berkeley course developed by Stuart Russell and Peter Norvig which was already used in 1200 institutions (http://aima.cs.berkeley.edu/). I used the textbook myself in my introductory AI class.

Sebastian Thrun, a roboticist without any prior involvement in educational technology, was also known as the leader of the team that won the DARPA self-driving car challenge, and, together with Peter Norvig, he was increasingly involved with Google. But after this powerful kickstart, the ball got rolling and soon courses which did not have already a prior reputation or which were not proposed by a university with a strong brand like Stanford became available and attracted an important audience. Time magazine declared 2012 "the year of the MOOC" and interest skyrocketed. By 2014 more than 6 million students had registered for a MOOC course.

Based on the unusual success of this educational format, Sebastian Thrun and colleagues founded a start-up company in 2012 called Udacity, soon followed by other start-ups, such as Coursera, also initiated by two Stanford AI professors, edX, set up as a non-profit consortium by Harvard and MIT, and Khan Academy, created by Salman Khan. These companies targeted university-level education. Quickly other companies such as Udemy or Canvas, sprung up that were teaching any subject (from yoga to guitar playing) and accepted courses by anyone who wanted to be an instructor.

Increasingly MOOC providers began to ask payment, with income shared between the instructor and the company, and some success stories started to circulate that money could be made by creating a successful MOOC. Meanwhile MOOC aggregators appeared, such as, http://www.curricu.me/. They survey and provide links to MOOCs that have been developed by different individuals and institutions. Also general platforms for building MOOCs were launched, such as http://mooc.org/, which is a cooperation between Google and EdX initiated at the end of 2013.

The well-established European long-distance organisations were somewhat caught by surprise. They had been delivering high quality course materials but always to smaller groups and therefore without ever reaching large audiences. Nevertheless, the European open universities are the natural competence centers for creating MOOCs in Europe. And indeed the UK Open University launched futurelearn.com in october 2013 and other European open universities banded together to create http://www.openuped.eu/.

A few start-ups for delivering MOOCs have sprung up in Europe as well, such as iversity (in Germany). Nevertheless, many European universities (such as Edinburgh university or Delft university) are delivering their MOOCs through American platforms (primarily Coursera) and it is getting more and more difficult for European platforms to compete, given the much bigger resources available to American providers thanks to venture capital. This is disappointing because Europe has a long tradition in distance education and engaged in a lot of research towards online courseware in the eighties and nineties.

4.2. Beyond the hype

In 2012, the first commercial MOOC providers made big claims that they were going to solve the crisis in education. Sebastian Thrun of Audacity proclaimed in Wired magazine that only 10 institutions would be needed to create all the educational materials required for the whole of the United States and that most colleges and universities would seize functioning within the next 50 years. Coursera anounced that they would provide Ivy-league level education online for free and that this would not only solve the escalating cost of college education in the US but also allow developing nations to close the knowledge gap. Some of the critical issues for scaling up education to sustain hundreds of thousands of course participants seemed to be highly underestimated. For example, Daphne Koller (CEO of Coursera) claimed that "*With some effort in technology development, our ability to check answers for many types of questions will get closer and closer to that of human graders.*" [12].

Exaggerated claims are a familiar feature of Silicon Valley start-up companies and they are accompanied by smart public relations that push these claims through the media. In this case the strategy worked again. Venture capitalists began to see a tremendous market opportunity, based on the belief that MOOCs might come to dominate education the way other companies such as Google have managed to dominate world-wide information access and thus control significant chunks of electronic commerce as a consequence. Thus, Coursera already raised (end of 2014) 80 million Dollar in capital—leaving competitors including the timid initiatives in Europe far behind in terms of available resources. This level of investment provides resources for speedy development so that some of the claims for MOOCs, which at first look exaggerated, could actually be made true.

On the other hand, the claims also generated a serious backlash, leading to an 'existential crisis' for (commercial) MOOCs by early 2014. University administrations saw a great opportunity for cutting cost, and hence possibly lowering tuition fees, by replacing or getting rid of their own courses and faculty. Some state governments began to work on legislation so that MOOC providers would be able to give official course credits, the same way existing private or public universities could. It was therefore to be expected that the faculty in colleges and universities that had been declared to become soon obsolete rose up against this development. One widely publicised case involved the San Jose State university philosophy department which refused to accept MOOCs as substitute for some of their own courses and challenged the (Harvard) course professors who had helped to create the MOOCs [9]. Partly as a result of these protest movements, efforts for accreditation of MOOCs in California were quietly dropped after a lot of debate.

What are the criticisms laying bare some of the serious limits of MOOCs in their current state of development and deployment? Are some of these limits inherent to MOOCs? Here are a few issues that have widely been debated in the blogosphere:

1. **Drop-out rate** Although the number of students that subscribe to MOOCs is huge, the drop-out rate is also very high (usually 90 %), which is much higher than any degree program in the regular higher education system. In defense of MOOCs, we have to say that there is no cost and no barrier to subscription and hence it is just as easy to give up. Many subscribers are simply curious. They want to see what the course is like, hope to find the time to study the subject, underestimate the difficulty, find a better course for their needs, etc. It is still amazing that tens of thousands do finish some of the high level engineering courses.

2. **Assessment and Grading** The most obvious critical problem for MOOCs is how to assess students. MOOC providers propose two vehicles: peer-to-peer grading and assessment through computer programs. For peer-to-peer grading, there is clearly a major problem, partly because students most of the time do not have the knowledge and experience themselves to come up with a decent evaluation of the work of others and partly because it is extremely difficult to get standardised grades across all students taking the course. Assessment through computer programs works for the basics in technical and scientific domains where the answers to tests are objectively known and unequivocal, but it is a very different matter when moving to human-related subjects such as history, music, writing, etc., or when creative answers are called for.

And how can cheating be avoided? Plagiarism is already rampant in the regular education system and technologies have sprung up to detect the unreferenced appropiation of texts. But in the case of distance learning, there is no way to be sure that the student who claims to follow a course actually does all the assignments him or herself. This is less of an issue in the case of continuing education where one might assume that the student is motivated to follow the course for his or her own benefit and certified credit is not the main objective, but it is a real issue for basic higher education (18-22 year olds) who are often mainly interested to pass a course with the least effort.

3. **Contextualisation** The claim that the MOOCs developed in the United States can simply be transferred to the rest of the world has also been criticised. Again, for very basic subjects, such as algebra or programming, it might be possible to come up with a 'universal' course. But for most other subjects, adaptation to the local context is one of the crucial ingredients to make a course relevant and successful to a particular group of students.

4. **Business model** At the moment MOOCs are mostly free. But many argue that this is bound to change as soon as there is a captive market. The typical modus operandi

of Silicon Valley style companies is to provide at first services for free to gain market share and wipe out possible competitors, but then at some point the balance tilts. We can expect the same pattern as with music, books, videoclips or other media content. At first these materials were also provided for free (possibly pirated) but once everybody was equiped and accustomed to electronic distribution and the traditional delivery channels (music stores and vinyl or CD records, printed books and bookstores, cinemas and DVDs) had collapsed, payment became obligatory either through advertising or direct payment. Those companies that had managed to win this race found themselves in a quasi-monopoly position and could ask the price they wanted or bully content providers. The history of Amazon is a case in point. Amazon managed to get a near monopoly on the electronic delivery of books, leading to a demise of book stories, decreasing income for publishers (and therefore less risk-taking) and a drop of income for most authors. This does not mean that electronic publishing is a bad idea, but somehow the creation of monopolies has to be avoided. For education this is even more true. It would be a tragedy for the richness of intellectual life if only 10 universities would remain in the US and, presumably, only a handful in Europe as well.

Given the cost of developing course materials (effective MOOCs require investments of several 100 000 EUR), it will not be straightforward to recuperate the investment and MOOC companies are already trying to find ways to monetise, such as by selling data about students to companies, by advertising textbooks, by acting as recruitment agencies bringing the best students in contact with companies seeking new applicants, etc.

Not all MOOCs are privately produced. Public universities and university consortia see the creation of MOOCs as part of their public service because it helps them to contribute to education as a public good. But that implies of course that the universities can continue to exist, and that MOOCs are not seen as a way to slash the resources available for higher education.

5. The Future of MOOCs

Today MOOCs are both praised as the future of education and heavily criticised as potentially having a disastrous effect on the quality of education and the commons of knowledge that is supporting it. But I believe we should not throw away the baby with the bathwater too quickly. On the positive side, the sudden rise and world-wide spreading of MOOCs is giving a tremendous boost to educational innovation and online distance learning. Distance education was already well under way in many countries, including in Europe. But it was seen as marginal and never reached the same large world-wide audience as MOOCs. This has clearly changed and so there are now resources leading to the development of exciting, more effective online learning environments [30] and tools for building them [7].

I believe that the following are realistic targets we can expect from MOOCs in the future:

- MOOCs are so far mostly used for continuing education. But this is nevertheless of the greatest importance. The time knowledge acquired in formal education remains valid is getting shorter and shorter, except perhaps for the very basic skills such as mathematics or writing. So everyone needs educational resources to keep up and expand their skills and MOOCs are ideally suited for this.

- MOOCs appear an excellent replacement for the classical textbook, even for textbooks that are online and are intended for self-study. It is therefore no surprise that libraries see a new role for themselves. For example, the New York Public Library organised in august 2014 the viewing of a MOOC with additional in-person help.
- MOOCs support blended learning. Educators can use the materials provided by MOOCs in their own courses and thus get more time and resources to emphasise the interactive and social aspects of education. An extreme form is known as the *flipped classroom*, where the lectures normally given by the instructor are replaced by a MOOC and the class itself is focused on motivating students, discussing the lecture material, doing additional exercises, social interaction between the students, and assessment and grading [27].
- MOOCs stimulate educators to create (more) attractive courses, simply because examples of excellent courses are now easily accessible. The demands of students also go up because they see what a good course look like.
- Search engines were in the past decade the 'killer app' for Artificial Intelligence technologies, specifically machine learning and natural language processing. MOOCs can play a similar role for the many valuable technologies that have been experimented with by AI researchers and educational technologists in the past decades, particularly technologies for developing student models and thus very accurately assessing the level of competence of students, integrating automated sophisticated problem solvers to customise challenges posed to students, and allow active adapted strategic planning how course materials are presented to the individual student.
- There are many domains that are currently falling outside the realm of 'official' education channels in many countries, particularly arts education. Many students who are eager to learn how to play instruments or participate in orchestras or bands are pushed towards private tutoring which their parents often cannot afford. Here MOOCs can have a major beneficial impact by giving access to novel learning resources [13]. The many technical challenges that need to be dealt with in order to achieve this properly are already pushing the technological development of MOOCs in exciting new directions [30].

One of the biggest challenges for MOOC developers is to escape the paradigm of instructional teaching, which is relatively easy to implement, and move towards the open-ended learning environments that are much more appropriate for a wide range of subjects from music learning to learning how to think and write. This implies going back to the original concept of cMOOCs [20]. I believe this can happen by striving to achieve *social MOOCs*: MOOCs that foster the self-organization of small communities of learners that collaborate over the internet and are supported by novel resources: learning materials to be explored in an open-ended way, peer assessment and feedback, automated feedback, easily navigatable links into information resources, a.o. Several papers in the present volume show very significant progress in this direction.

But in the end, education is more than acquiring practical skills or knowledge. Particularly the 18-22 year old segment of the population, which is the primary population for higher eduction, is not only acquiring specific skills and competence. These students are still in the process of learning about life, how to work, how to get organised and live harmoniously with others. Teachers play an important role in this process, picking up where parents left off, and a lot of learning is peer-to-peer. The social relationships that are established in college settings are the first really useful social network for succeeding later in life. This task cannot be done at a distance by MOOCs. It requires being together, engage at a human level, and getting access to educators with a rich life experience.

From this perspective, a recently started experiment by Black Mountain College in North California is intriguing. Black Mountain College was already famous in the nineteen forties for its innovative approach to education with teachers such as Buckminster Fuller and John Cage. Recently the college started a MOOC campus (http://mooccampus.org/) which does completely away with the regular teaching by faculty and replaces that aspect of education with MOOCs. It argues that "*many of the best lessons from school come from the social aspects, relationships, connections, and extracurricular activities that have traditionally been part of the college experience.*" And so the school brings together motivated individuals which do not necessarily want a diploma (although degree-seeking students are welcome as well) but want a DIY (*Do It Yourself*) education, motivated and supported by like-minded peers. The role of faculty is similar to a team of coaches. They provide resources and organise to some extend what MOOCs have to be studied, but they focus mostly on the personal growth of the individual students and on the stimulation of a healthy social climate within the group. Is this a possible future of education?

6. Acknowledgement

This research was supported by the FP7 Technology Enhanced Learning Program Project: Practice and Performance Analysis Inspiring Social Education (PRAISE) and by an ICREA research fellowship from the Catalan government supporting the author at the Institut de Biologia Evolutiva (UPF-CSIC) Barcelona.

References

[1] Beuls, K. and J. Loeckx (2015) Steps towards Intelligent MOOCs. Chapter 9. In: Steels, L. (ed.) (2015) Music Learning with Massive Open Online Courses MOOCs. IOS Press, Amsterdam.

[2] Breslows, L. et al. (2014) Studying learning in the Worldwide Classroom. Research into edX's first MOOC. Research and practice in assessment. Vol 8. pp. 13-25.

[3] Eisenstadt M. and T. Vincent (2000) The Knowledge Web. Learning and Collaborating on the Net - Knowledge Media Institute, Open University, Milton Keynes, UK.

[4] Eurostat (2012) Key data on education in Europe 2012. Education, Audiovisual and Culture Executive Agency of the European Commission, P9. Eurydice.

[5] Garrouste, C. (2010) 100 Years of Educational Reforms in Europe: a contextual database. Joint Research Centre, European Commision, Ispra.

[6] Giol, F. (ed.) (2009) Lectures contemporaines de la crise de l'èducation. L'Harmattan, Paris.

[7] Guo, P., J. Kim and R. Rubin. (2014) How Video Production Affects Student Engagement: An Empirical Study of MOOC Videos. In: Proceedings ACM User Interface Software and Technology Symposium. Honolulu.

[8] Haber, J. (2014) MOOCS. The MIT Press, Cambridge Ma.

[9] Hadreas, P. et al. (2013) An open letter to Professor Michael Sandel from the philosophy department at San Jose State University. http://chronicle.com/article/The-Document-an-Open-Letter/138937/

[10] Illich, Ivan (1971). Deschooling Society. Harper and Row. New York.

[11] Kahn, K. (2015) A half century perspective on the role of computers in learning and teaching. This volume.

[12] Koller, S. (2011) Death knell for the lecture: Technology as a passport to personalized education. The New York Times December 5.

[13] Loeckx, J. (2015) Learning Music Online. Chapter 2. In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

[14] Okada, A., Buckingham Shum, S., Bachler, M., Tomadaki, E., Scott, P., Little, A. and Eisenstadt, M. (2009) Knowledge media tools to foster social learning, in eds. S. Hatzipanagos and S. Warburton (eds.) Handbook of Research on Social Software and Developing Community Ontologies, IGI Global, London. pp. 357-380.

[15] Papert, S. (1994) The children's machine: Rethinking school in the age of the computer. Basic Books, New York.

[16] Quere, M (2012) L'Etat de lécole. French Ministry of Education, Paris.

[17] Self, J. (1999) The defining characteristics of intelligent tutoring systems research: ITSs care, precisely. International Journal of Artificial Intelligence in Education (1999), 10, 350-364.

[18] Schnepf, S. (2014) Do tertiary dropout students really not succeed in European labor markets. Discussion Paper 8015. IZA, Bonn.

[19] Selingo, J. (2014) MOOC U: Who is getting the most out of online education and why. Simon and Schuster, New York.

[20] Siemens, G. (2005) Connectivism: A Learning Theory for the Digital Age. International Journal of Instructional Technology and Distance Learning, 2(1).

[21] Smith, P. (ed.) (2003) Violence in schools. The response in Europe. Routledge, London.

[22] Smith, S., B. Sherwood (1976). Educational Uses of the PLATO Computer System. Science 192 (4237). pp. 344-52.

[23] Steels, L. (1992) Kennissystemen. Addison-Wesley, Amsterdam.

[24] Tokoro, M. and L. Steels (eds) (2004) A Learning Zone of One's Own, IOS Press. Amsterdam.

[25] Tokoro, M. and L. Steels (eds) (2005) The Future of Learning. Issues and Prospects. IOS Press, Amsterdam.

[26] Vallima, J. and D. Hoffman (2008). Knowledge society discourse and higher education. Higher Education, 56(3), 265-285.

[27] Van der Perre, G. (2015) Blended learning and MOOCs. Chapter 15. In: Steels, L. (ed.) (2015) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

[28] Van den Branden, J. and G. Van der Perre (1997) EuroPACE 2000. European Journal of Education. 32(4). pp. 359-367.

[29] Whitfield, D. (2000) The third way for education: privatisation and marketisation. FORUM, 42(2), pp. 82-85.

[30] Yee-King, M., M. Krivenski, H. Brenton, A. Grimalt-Reynes, and M. d'Inverno (2015) Music Circle: Designing educational social machines for effective feedback. Chapter 7. In: Steels, L. (ed.) (2015) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

[31] Young, J. (2013) Beyond the MOOC hype: A guide to education's high-tech Disruption. Chronicle of Higher Education. Washington.

[32] Zajda, J. (ed.) (2006) Decentralisation and privatisation in education: the role of the state. Springer-Verlag, Berlin.

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-21

Chapter 2. Learning Music Online

Johan Loeckx
VUB Artificial Intelligence Lab, Brussels, Belgium

Abstract. Much of music learning happens outside the classroom in an informal setting, for example online by watching YouTube movies. Despite the many research effort and government projects, many online lack pedagogical foundations while few music educators fully employ online technology yet. In this paper, existing research on online music learning is reviewed as well as platforms and sites that are currently available, mapped in a so-called "market quadrant". Also, the main problems with the current state-of-the-art are identified, along promising technologies that are believed to address these challenges.

Keywords. music learning, online learning, adaptive learning, intelligent tutoring, MOOCs

1. Introduction

It is a running joke among musicians to fret that learning music requires more discipline than joining the army. Mastering an instrument takes upto fifteen years of hard study and truly understanding music is never-ending effort. Even Beethoven took lessons in counterpoint when he was already composing. Maintaining motivation over the long term as well as exquisite personalized feedback are key to success. In principle, online learning thus seems like a perfect candidate to facilitate and increase the intensity of music learning.

Online learning in general is on the rise for more than a decade now [1] and has been found to be at least as effective as traditional teaching [34]. This effectiveness has been mainly attributed to pedagogical aspects rather than the delivery medium itself, however. It appeared that online learning was most effective when the medium and technology was used *comprehensively* in support of learning instead of its usage being a goal on its own. Despite all this, online learning is still in an embryonic phase, evidenced by the fact that most online courses are still reflections of traditional face-to-face teaching. This is a typical pattern in technology adoption in education, where technology is first used to improve what we are already doing (replication), then to do things that we could not do before (innovation), and finally engaging technology to transform content, pedagogy and learning in a fundamental way (transformation). This is even more so for the online learning of music, where adoption is scant and research on the effectiveness is quasi unexisting [52]. At the same time, a lot of groundbreaking work has already been performed in the field of distance education and pedagogy in general, that could significantly improve online music teaching.

This article surveys the current state-of-the-art in online music learning in an effort to guide teachers and platform designers and argue for the need to bring pedagogy back into the equation. Existing commercial and experimental sites and platforms will be summarized in a so-called *market quadrant*, and current pedagogical issues will be examined. Finally, we will glimpse how future technology might solve some of these issues and provide unique opportunities for scalable online learning.

2. What is already known...

Online learning is a young field that has evolved organically from distance learning, with the internet just offering a new kind of delivery mode. Though a large body of research exists on distance learning, online learning is still relatively underinvestigated. This is particularly the case for learning music online. There have been many research efforts and government projects but, despite the eminence of technology in our everyday's life and educational discussions, in general, music educators do not yet extensively utilize technology directly with students to improve learning [4]. The field has been slow to adapt new technologies due to a combination of skepticism and unfamiliarity [42]. Experiments are set up primarily to test feasibility, with little attention to learning effectiveness and rigorous validation.

It is striking that in times that complete courses and degrees are being offered online, aspects like curriculum design and pedagogy—central to create high-quality learning experiences—are being largely neglected [34] and little research is being done on the effectiveness of these programs [52]. This is even more remarkable given the long, well-documented history of distance learning, dating back more than 200 years [55]. Excellent research has been performed, for example, in the context of Open Universities—such as the work by Rees in summarizing distance learning in music education [42]. For this reason, in the next sections, a brief summary is given of what is known about online teaching. We will focus on three aspects that make online learning unique and different from more traditional forms of teaching: (1) the central role of technology, (2) the distance between teachers and students and (3) the opportunities for (online) collaborative learning.

2.1. Teaching with technology

As the online context increases the complexity of interactions between content, pedagogy and technology, these different aspects can and may not be considered independently: knowing *how to use* technology is not the same thing as knowing *how to teach* with technology [36]. Yet, analysis of existing research has revealed that studies tend to choose technologies, content and learning activities arbitrarily rather than following a theory [34].

Too often, however, online courses are the digital counterparts of traditional face-to-face lectures. This is even more so in the domain of music, where music instructors are often unfamiliar with technology or even resistant to it. Yet, "quality online education will be realized only when traditional views of content and pedagogy are reconceptualized within new frameworks that include technology" [49]. Online courses thus need to be designed specifically with the intended delivery platform in mind. Only this way it can be ensured to both make optimal use of the technology at hand at the service of *pedagogical purposes*, and take into consideration the various constraints and specificities of the learning environment [8].

In an effort to create a holistic vision on teaching with technology, the TPACK framework [36] has been devised to include pedagogy, knowledge and technology. Roughly speaking, TPACK is a framework for integrating technology into instruction, arguing for a unified approach in which subject matter and technology are developed in companion. Rather than considering them separate bodies of knowledge, the framework stresses the *complex interaction between them*, needed to create high-quality effective learning experiences. Traditionally, teachers need to both have a deep understanding of the subject matter (content knowledge, CK) and possess the necessary pedagogical knowledge (PK): knowing how to teach, to motivate students, test their skills, scaffold concepts, etc. We believe that for technology to be a useful tool for teaching and learning, teachers should have a good grasp of the technology itself too, as well as how it interacts with content and pedagogy. So far, there has been little interest yet to incorporate these kinds of frameworks into the context of online learning [49].

TPACK has been extended for music instruction by Bauer [4,5]. The music learning activities have been designed to expand an instructor's existing array of practices to include technology—making it an excellent start for teachers who want to increase the use of technology in their music courses. Rather than focusing on the technology itself, the activities attempt to bring attention to the *learning action* students perform when they carry out the activity [19]. Other frameworks can often be adapted or interpreted in the context of online learning too. It would lead us too far to go through all the existing frameworks; we refer to excellent review articles on the topic [33,12].

2.2. Minimizing transactional distance

Another prominent feature of online learning is the *distance* between teachers and students. Already in 1972, M. G. Moore observed that distance was not simply a geographical separation, but a pedagogical concept indicating the *transactional distance*, i.e. the resulting psychological and communicating space that needs to be bridged [38]. In his visionary paper *Learner autonomy: The second dimension of independent learning*, Moore presented an optimistic outlook for distance learning, one in which idiosyncrasy and creativity should be cherished and nurtured for every personality to mature to the level of self-actualisation. According to Moore, three elements are key to determine the extent of transactional distance:

- **structure**
  Structure is needed in order to guide students to acquaint themselves with the different concepts and aspects in a course. Structure includes the course

design, scaffolding strategy, the organisation of learning experiences and the use of various communication media.

- **dialogue**
  Synchronous or asynchronous interaction between teachers and students and students mutually. Though the *medium* of communication is central in the design of dialogue, it is not sufficient. A discussion board, for example, can enable two-way interaction but does not guarantee that it occurs or that it decreases the transactional distance accordingly.
- **learner autonomy**
  Personal involvement of learners drastically improves teaching effectiveness. The expression of the learner's personalities, and time for individual reflection are needed to improve self-regulation. Both meta-cognitive (competence in estimating challenge levels and awareness and monitoring of own capacities) and non-cognitive skills (character, emotions, creativity) are examples of this category.

Often a balance should be found between the different aspects. For example, the amount of content covered should be balanced with the amount of time required for reflection and discussion.

Neuroscientific research supports this viewpoint, indicating that the more ways learners interact with a concept, the more secure it becomes in memory. Furthermore, people tend to learn more efficiently if what they learn is personally meaningful, as well through social interactions. These aspects become even more essential in an online setting where the instructor is not physically present and thus can not actively and continuously gauge the student's interactions and understanding to adapt the teaching style accordingly. All these findings provide support for the four dimensions of transactional distances have been identified in [21], centred around the interactions between:

- learner–content,
- learner–instructor,
- learner–learner, and
- learner–interface.

A similar idea is formulated by the *Community of Inquiry* framework that claims that effective learning occurs in a *community of teachers and students* through the interaction of social, teaching and cognitive *presences*, an identity created through interpersonal communication. A community of inquiry is defined as a group *who collaboratively engages in purposeful discourse and reflection to construct personal meaning and confirm mutual understanding* [17]. Figure 1 shows an overview of the different kinds of interactions.

2.3. Collaborative learning

One of these dimensions, learner/learner interaction, has become particularly popular in the last years with the success of social media and the coming of MOOCs. It is seen as one of the holy grails to solve the scaling problem in MOOCs, together with adaptive learning [58]. This learner/learner kind of learning is often referred to with different names like peer learning, or *collaborative learning*. Exist-

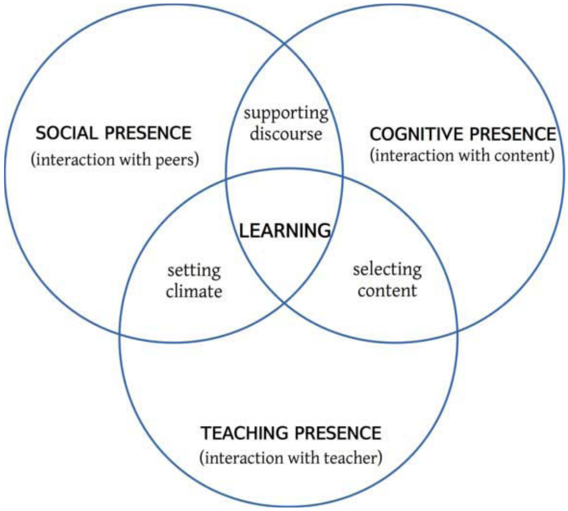

Figure 1. Interactions between different kinds of presences in the community of inquiry framework. Effective learning occurs in a community when social, cognitive and teaching presences interact. In an online setting too, it is important that all interactions are taken care of.

ing definitions of collaborative learning differ significantly and are often vague or subject to interpretation. Roughly speaking, collaborative learning is a situation in which two or more people learn through interactions [15]. This means that collaborative learning can not be reduced to one single mechanism: just like people do not learn because of their individuality, but rather because the activities they perform trigger learning mechanisms, people don't just learn collaboratively because they are together. Rather, it are the interactions between the create activities (explanations, mutual regulations,...) that trigger cognitive learning mechanisms (elicitation, internalisation, ...). For this reason, facilitation and support for these activities should be taken into consideration when designing collaborative learning experiences.

3. What already exists...

Having discussed the main unique pedagogical characteristics of online learning, we will now give an overview of the technological state-of-the-art. Beside collaborative learning, technology is seen as the other promising road towards scalable learning [58] and due to the popularity of Web 2.0, mobile devices and social media, the promise of online learning has become within reach. The nexus of ubiquitous connectivity, cloud computing and intelligent software opens up the opportunity to reduce the transactional distance (1) by boosting the interactivity between learner and content and (2) by providing a virtual "substitute" for the instructor. The following summary will not be limited to formal learning only, as a lot of students' musical achievements occur outside (music) schools. Moreover, with online social networks booming and becoming more pervasive, the borders between formal and informal learning are fading and a lot of musical achievements take place online [44].

3.1. Market quadrant

A quick search on the web will return an abundant amount of sites that claim to “make music learning simple”. Clearly, this is only marketing talk as any musician can tell that learning music is a hard process that takes many years, genuine motivation and loads of discipline. Close guidance over the long term to construct a scaffold of personalized learning experiences is thus indispensable. For this reason, two aspects will be covered in deeper detail: the level of *curriculum design* and the level of *personalized learning.*

To gain oversight, the coming overview is mapped out in a so-called “market quadrant”, shown in Fig. 2, that maps out every solution with respect to two axes[1]. Curriculum design refers to the extent that the learner is guided through the material, using instructional techniques like scaffolding and indications of what is next. Personalized or Adaptive learning is a term for a collection of techniques—human and computerized—to provide a unique and personalized learning experience, aligned with the goals and needs of each student.

These abstract concepts above have been translated to quantifiable indicators for each axis, based on pedagogical theory. The following components are commonly considered when defining curriculum [26,6]:

(a) *subject matter*: the scope, integration and depth of the material discussed;
(b) *instructional plan*: the strategies of organizing and dividing content;
(c) *horizontal organisation*: the order and continuity of content;
(d) *assessment and evaluation*: the criteria for examining the results.

The indicators listed above and below have been operationalized and have been explained in some more detail in tables 1 and 2 in the Appendix. To measure the level of personalized learning, the following aspects have been used [57]:

(a) *student modelling*, or the ability to represent a student’s knowledge;
(b) *expert modelling*, the representation of the expert’s knowledge ;
(c) *instructional modelling*, the adaptation of the teaching to the student;
(d) *generativity*, the ability to generate appropriate problems for the student’s need;
(e) *mixed initiative*, the interaction with a student to respond usefully; and
(f) *interactive learning*, the learning activities that require real-time student engagement in a domain-relevant and contextualized manner.

When online music technologies are plotted in a graph, we can roughly identify 4 *quadrants.* (a) *instructional sites* with a fixed set of web pages disclosing a particular topic; (b) *games & training tools* that focus on a small particular aspect of music learning; (c) *social media* that empower peer interaction to achieve personalized learning; and finally (d) *blended models*, experimenting with new pedagogies and technologies. A quick glance at the market quadrant in Fig. 2 shows that there is a so-called ”scalability” gap that needs to be bridge to achieve true scalable learning online. In the following sections, we will discuss each of these quadrants in more detail. References to each of the different platforms can be found in the Appendix.

[1] The source data of the graph can be found in the Appendix in Tables 3, 4 and 5

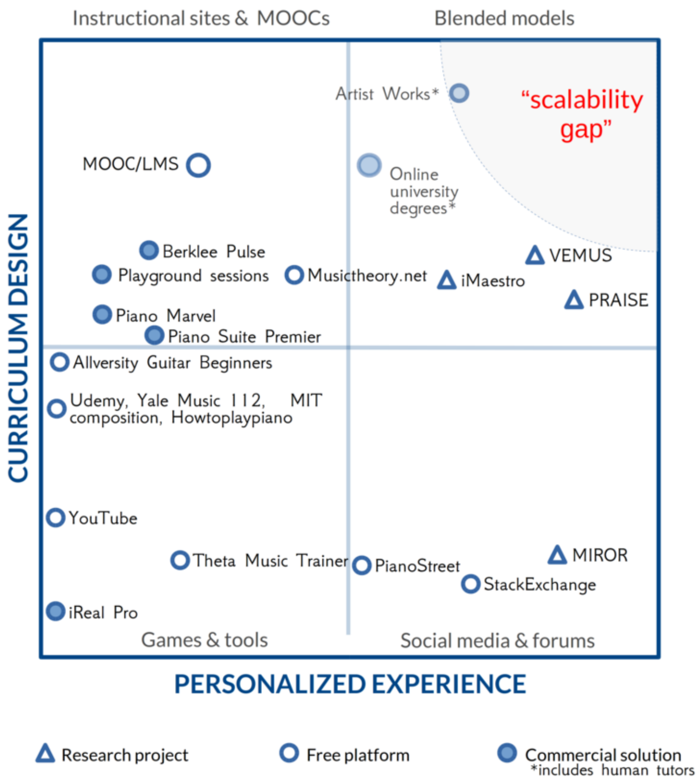

Figure 2. Learning music requires a personalized approach and motivation over a long term (5–15 years). Careful curriculum design and the fostering of personalized experiences is therefore crucial to successful online learning. This market quadrant plots current solutions with respect to these two axes. A "scalability" gap is observed, currently hampering the true potential of online scalable learning.

3.2. Instructional sites

A first category represents the largest share of online music learning resources. Emerged from traditional websites, they consist of typically high quality resources in the form of a combination of text, videos, images and schematics to explain music theory, composition, harmony, etc... Udemy, Yale, MIT, merlot.org and howtoplaypiano.ac are representative examples of this category.

Typically, the content is organized in a fixed curriculum, or a (tagged) collection of independent objects. Online music degrees offered by higher education institutions and delivered through Desire2Learn or other Learning Management Systems, are part of this category as well. Though they offer high quality content, these type of sites often lack the personalised and interactive experience needed for mastering music, especially an instrument. Also, their passive nature is barely inspiring or motivating. The coming of MOOCs like Coursera, however, has opened up new possibilities by combining the advantages of traditional curricula with feedback from peers. Current courses on these platforms, however, show the same shortcomings and lack multi-dimensional assessment.

Nevertheless, some very interesting semi-interactive tutorials exist that carefully explain the basic concepts of music theory, (like Musictheory.net), and in rare cases also advanced topics like Counterpoint or Fugue analysis (Merlot). Proprietary commercial solutions like Berklee Pulse, Piano Marvel, Piano Suite Premier and Playground sessions, offer a more "musical" approach but lack quality feedback, assessment and rigorous curriculum design. Basically they consist of instructional videos, enriched with a song database that allows students to play along.

3.3. Specific games & training tools

A second category consists of interactive games and training tools that promise easy learning of music. Typically they focus on very specific musical problems and thus exhibit very little curriculum design maturity and provide no clue to "what's next". Examples of these tools are: ear training, interval training, playing the right note, tuners, chord finders and playing/singing along. Theta Music Trainer offers an interesting collection of these online tools. Though these tools are designed for entertainment ("fun") and are highly interactive, they consist of disconnected pieces of training lacking scaffolding and offer only superficial feedback like right/wrong, too late/early, ... Though some tools offer a manual setting of difficulty level, these tools are not personalized and lack a pedagogical basis in musical development theory. Their limited scope restricts their effectiveness too.

3.4. Social media and forums

Research has shown that learning can occur through shared activities performed in participatory communities of practice [53]. Self-directed learners that are in control of their learning, can construct environments to support their learning (for example in social networks). Equivalent processes in the real world of music are garage bands, choirs and orchestras. A reflection in the online atmosphere makes up the third and relatively new phenomenon in music learning, the emergence of social media or networks. Too often, however, these online activities are seen as asocial by traditional educators [46], probably because this idea is at odds with traditional view on (music) education that focuses on the transfer of factual and established knowledge from a central authority.

Examples of platforms on which musicians can ask specific questions are Stack Exchange and PianoStreet. Though these Question and Answering platforms offer highly personalized and collaborative feedback and interaction, the ad hoc "requirement-based" approach offers little guidance and combined with the small scope of questions, again limits their effectiveness for music learning in the long term. Still, the impact of general platforms like YouTube, Facebook and Soundcloud should not be underestimated in informal settings, creating communities of musicians that share music, get feedback on compositions and engage in discussions.

3.5. Blended models

The last kind of platforms, *blended models*, finally, contains sites that try to combine a rigorous curriculum with personalized learning. Three main techniques can be discriminated. The first blends learning settings and human-assisted tutoring with curriculum design. Online university degrees are good examples in which a purely online degree is enhanced by feedback from human experts through synchronous or asynchronous feedback. Another example is the solution from *Artist Works*, in which students and human tutors communicate by uploading video of their playing and get feedback through another video. This way, students receive personalized feedback from experts, who are also in charge of mapping out the curriculum. Clearly, this solution does not scale.

A second approach attempts to improve personalized learning by putting peer learning and communities at the core of traditional instructional sites like MOOCs. The discussion boards are no longer secondary to the learning material, but central to the learning experience. cMOOCs are manifestations of this category. A similar approach is taken in PRAISE, a European FP7 project. PRAISE (Practice and peRformance Analysis Inspiring Social Education) focuses on collaborative online processes and aims to analyse the nature of feedback that characterize those processes [47].

Recently, *adaptive learning* is hoped to replace some tasks of a human (expert/peer) tutor by intelligently designed pieces of software that guide students through the learning process; this is done by selecting appropriate learning activities and providing interactive tutoring and sensible feedback. The Virtual European Music School's (VEMUS) is such an example. It tries to pave the way to an open interactive platform [11] by integrating a whole range of technology, connected to music learning: interactive on-line and off-line collaboration "music rooms", automated audio analyses, score editing tools etc.. The iMaestro project [23] took a more technological approach and has focused on creating tools and methods for courseware production and to host interactive, collaborative and creative rehearsals, incorporating multi-modal interfaces and new kinds of visualizations. Both have a very technological focus and lack in-depth, intelligent and musical feedback, however.

A last research effort is worth mentioning, as it takes a fundamentally different approach to music learning. The MIROR project is based on so-called "reflexive pedagogy" [35] or the idea that children can learn how to play music by listening to virtual copies of themselves. Specific machine learning software have been developed, called *Interactive Reflexive Musical Systems* (IRMS). The system acts as a tutor, designed to interact with learners in the field of music improvisation. Several artificial intelligence and multi-modal signal processing techniques were employed to create an inspiring environment that fosters creativity. The evidence on learning effectiveness remains meagre, though.

4. What the problems are...

So far, we have outlined the existing pedagogical knowledge on online learning, as well as the state-of-the-art in educational technology for music learning. As mentioned before, there are, however, still many challenges that the platforms above are facing. In the next sections, we will dig deeper into these issues.

4.1. Limited body of research

Despite the large body of research on *traditional* music pedagogy, most of the current platforms lack a rigorous pedagogical basis and are primarily inspired on a behaviourist / transmission model view on education. Existing solutions are typically constructed in an ad-hoc fashion and lack systematic testing based on theory [34]. Even worse, there is very little scientific evidence on the effectiveness of these programs at all.

Although growing, the amount of research on *online* music learning is still very limited [8], and the studies that do exist are typically small in scope: they consider a small number of students (10-30) and experiments last a short period of time only (few weeks to few months). Furthermore, most online degrees deal primarily with academically oriented courses like music theory or history, not the teaching of manual skills [41]. For example, the different pedagogical functions of synchronous vs. asynchronous communication have been largely uncovered, even though existing research has shown that asynchronous communication when assessing or discussing music performance has been shown feasible and functional but less effective than synchronous face-to-face interaction [14]. As music is intrinsically multi-modal, more research is desperately needed for reliable multi-modal interfaces and feedback, even though technology is improving.

Finally, as a lot of these studies have been carried out by technicians and computer scientists, they typically focus on technology and delivery mode, neglecting pedagogical aspects and missing the crux of how the new technologies can be put to effective use [8].

4.2. Lack of quality assurance

Second, for online learning in general, a lot of critical voices are heard from professors in universities and colleges, afraid that "cheap MOOCs" with underperforming pedagogical foundations will replace faculties, warning of the misleading promise from MOOCs to improve education [37]. Exactly the opposite reasoning is heard too, however: that MOOCs may as well *improve* the quality of courses globally, as locally organized courses can no longer afford to be of a lower standard than an online course [2].

From the moment that universities or other institutions like governments are starting to provide certificates of equal value to online courses, the *credibility of assessment* becomes critical [10,18]. There are still, however, some compelling challenges like cheating or scaling to overcome [51]. Research into peer and self-assessment look promising and seems to work well when the students level is homogeneous, when students have a similar model of perception of "quality" and if they are properly trained in grading. Furthermore, evidence seem to indicate that students *learn* from these kinds of grading [45]. Still, assessments are in an early stage of development, testing for skills only superficially and not yet in a trustworthy way.

Furthermore, as mentioned before, very little attention is paid to assessment of the *learning effectiveness* of existing degrees and platforms. A study by Harvard in the context of MOOCs, revealed that students “emphasized the continuing importance of in-person discussion sessions” [20], as contact and dialogue between tutors and students is still considered key to teaching [29,54]. This necessity of human tutors makes some people fear that MOOCs will in the end not keep up to their promise, because maintaining the staff-to-student ratio of roughly 1:25 needed to counsel students will mean that even MOOCs will not scale. For this reason, Laurillard concludes that “education is not a mass customer industry: it is a personal client industry” [31].

4.3. Little personalized learning

The current standardized nature of online degrees and platforms, lacking personalized instructional guidance, thus makes current MOOCs suboptimal platforms. With the exception of some highly self-motivated and disciplined students, the lack of supervision and guidance hinders effective learning [39]. Also, current courses do not take into account personal differences in learning goals and ambitions, their background or the pace of learning. Especially with a diverse audience as in MOOCs, these factors are crucial for success [3]. In addition, learning material is typically presented in a linear “transmission model” fashion—not effectively exploiting pedagogies and other technologies suited for the online realm, such as those rooted in artificial intelligence like student modelling, gamification, intelligent tutoring systems and simulation.

5. ...and how technology can help

One of the main reasons why the huge potential of online courses is not exploited, is the tendency to think in terms of classical course-room lectures. The most interesting experiments have thus emerged by challenging the status quo by loosening the rigid control of learning outcomes and assessment criteria, shifting from content delivery to a learning process, experimenting with alternative assessment methods, games, and so forth [29]. Despite all the promises and expectations, the effects of computer aided instruction and virtual learning environments on learning have been rather disappointing till now.

The ultimate question is whether collaborative learning and novel intelligent software will be able to bridge the ”scalability gap” illustrated in 2, delivering a personalized learning experience and a rigorous curriculum to the masses. The omnipresence and advancement of information technology and the hype for MOOCs may have created a new momentum for online learning. Indeed, web technology has advanced to the stage that complex applications can run in the browser (e.g. Office 365), while digital devices and social networks like Facebook have become an integral part of our lives. Multi-modal interfaces like sound, video, or motion capture are becoming ever more sophisticated and open unique opportunities to provide a rich interaction with the user. In the following sections, we will elaborate on the unique opportunities that the current technological advances with a global leverage might bring.

5.1. Educational Data Mining

A first interesting evolution is the emergence of a field called ”Educational Data Mining” (EDM), that should be seen in the larger trend towards digitalization of education [56]. EDM aims to apply data mining techniques in order to better understand learning. Its applications range from analysis and visualization of data, providing feedback to instructors and course designers (when do students login, what topics do they visit the most, what patterns can be identified), making recommendations to students (for example of related content to explore), predicting student performance, student modelling, learning analytics and much more [43]. LearnSphere, an online database aimed at collecting information on learning, is another shot at consolidating the knowledge on learning and unlocking its potential to create better online courses [32]. Similarly, data scientists will have a seat at the table when designing courses, for example, to translate learning hypotheses into measurable quantities [28]. Studies start to arise that empower MOOC platforms to gauge the knowledge and competences of students and their capacity to learn [16].

5.2. Adaptive learning / intelligent tutoring

Another obvious shortcoming in music learning online, are personalized curricula that go beyond fixed-content delivery, difficulty levels (beginner/medium/ advanced), or very small-scope training tools. Domain models, as investigated in artificial intelligence's intelligent tutoring systems or recommender systems, should be constructed for the domain of music. Strangely enough, students have mainly remained out of the picture in existing online platforms. Keeping track of student's curriculum, emotions, practice schedule and their musical knowledge through student modelling is an absolute necessity. Feedback, tailored to the specific learning goals and problems a student encounters (diagnostics and repairs), could drastically improve the learner's experience [7]. Furthermore, the commonly used behaviourist approaches are not particularly well suited to foster creativity and critical thinking. Yet, creativity—diminishing since the '60s [30]—is considered as the most crucial factor for future success and innovation, according to a survey of IBM [22]. Though teamwork and collaboration can stimulate creativity [25], more attention should be paid; for example, using exploratory interactive learning environments.

5.3. Personalized feedback

Automated feedback provided by existing interactive learning environments currently lacks musicality. Knowing that a note is too early or late does not provide the necessary means to correct this situation. Furthermore, offering real help in the sense of providing learning strategies to improve is essential to speed up learning and make it inspiring. In addition, the automatically generated feedback on performances is highly unidimensional and lacks qualitative components that are crucial in learning music: rather than indicating wrong/right or note too early/too late, repair mechanisms should be suggested to improve their practice.

For beginning students, this kind of interaction does not motivate nor help them to advance. Also, instructional techniques like scaffolding are not employed.

One option that is seen as the "holy grail" of online education is to employ intelligent tutoring systems to guide students through the learning experience and provide them individualized sensible feedback. These systems should be designed to keep users in *flow* by presenting them challenges and experiences that match their capabilities at any time, prevent boredom or anxiety and optimize learning [13,48].

5.4. Gamification

A last promising approach to online learning is gamification. Teachers have been using games for ages in their classes: a recent survey unveils that 73% of the teachers in the USA use games at least once a week [24]. Gamification and "serious games"—the use of concepts and techniques from games outside the entertainment business—are currently explored [40]. One recent example is the MineCraft-EDU project, reaching about 250,000 students [50], where pupils can learn about history in a virtual world representing ancient societies, meet famous historical figures, embark on quests to learn about these fascinating cultures, explore genetics in real-time interactively [27], simulate social conflicts, or get insight into the spread of epidemics through a simulator game [9].

6. Conclusions

Learning music is slowly moving into the digital and online atmosphere. Interesting opportunities lie ahead, with the coming of multi-modal interfaces and global communities of learners. The advancements in artificial intelligence may provide personalized automated learning experiences, employing intelligent tutors and gamification techniques, while the field of educational data mining and learning analytics may provide key insight into how people learn, finally.

However, there are still some issues that need to be addressed to bridge the "scalability gap". First, creators of online courses should focus more on pedagogy and on how to teach with technology, reusing the large body of knowledge that has been constructed in the domain of distance learning. Second, further pedagogical study is needed about the techniques required for teaching manual skills using new interfaces. Lastly, more research is needed to assess and investigate the effectiveness of online learning platforms. For the domain of music specifically, more experimentation is needed to investigate efficient modes of synchronous and asynchronous communication—for performance practice and music learning in general.

A. Indicators & data

Aspect	Indicator
Subject matter	
scope	Number of handled subjects [2]
integration	Are the educational objects organized using tags or categories?
	Are their dependencies between educational objects?
Instructional plan	
scaffolding	Are learning tasks organized according to difficulty?
	Are tasks organized conforming a logical structuring of content?
Horizontal organization	
continuity	Is the learner guided through the subject matter?
sequence	Is there a strategy behind the *sequence* of learning tasks?
Assessment	
tests	Is there some form of assessment?
	Are the tests more sophisticated than multiple-choice?

Table 1. Criteria of maturity of Curriculum Design. Every studied software, site or platform is scored on each of these categories.

Aspect	Indicator
Student modelling	
skill level	Is the knowledge of the topic being tracked?
goals	Can the learning be customized?
Expert modelling	
domain model	Is there a model of the domain knowledge?
Mixed initiative – Interactive learning – generativity	
feedback	Does the feedback provide insight?
	Is qualitative feedback given? (typically by humans)
exercises	Can new exercises be generated?
game play	Are there games in which real-time interaction is required?
just-in-time	Is information given on the right time, situated an on-demand?
Instructional modelling	
skill level	Is the skill level adapted to the student's need?
repairs	Are problems diagnosed and repaired?

Table 2. Criteria of Adaptive or Personalized Learning, from literature on Intelligent Tutoring Systems.

Platform	URL	Type	X-score	Y-score
Coursera	https://www.coursera.org/	F	1.88	1.71
Musictheory.net	http://musictheory.net	F	2.19	2.82
Allversity Guitar Beginners	http://www.allversity.org/courses/beginners-guitar	F	0.50	3.49
StackExchange	http://music.stackexchange.com/	F	4.50	5.71
YouTube	http://www.youtube.com	F	0.50	5.04
Theta Music Trainer	http://trainer.thetamusic.com/	F	1.73	5.49
Udemy "Learn free music theory"	https://www.udemy.com/learn-free-music-theory/#/	F	0.50	3.93
Howtoplaypiano.ca	http://howtoplaypiano.ca/	F	0.50	3.71
Music 101: Intro to Music	http://education-portal.com/academy/course/	F	0.50	3.27
BBC Sing	http://www.bbc.co.uk/sing/learning/	F	0.50	4.38
Yale "Music 112: Listening to music"	http://oyc.yale.edu/music/musi-112	F	0.50	3.49
MIT Introduction to Music Composition	http://ocw.mit.edu/courses/music-and-theater-arts/	F	0.50	3.49
Online music theory tutor	http://www.childrensmusicworkshop.com/musictheory/	F	1.73	3.71
PianoStreet	http://www.pianostreet.com/	F	3.42	5.71
iReal Pro	http://irealpro.com/	C	0.50	6.16
Piano Suite Premier	http://www.adventus.com/store/piano-suite-premier/	C	1.42	3.27
Piano Marvel	http://www.pianomarvel.com/take-a-tour	C	0.96	3.04
Artist Works Music & Arts campus	http://artistworks.com/	C	4.65	1.27
Playground sessions	http://www.playgroundsessions.com/	C	0.96	2.82
Berklee Pulse	http://www.berkleepulse.net/	C	1.42	2.60
MIROR	http://www.mirorproject.eu/	R	5.27	6.16
PRAISE	http://www.iiia.csic.es/praise/	R	6.04	3.04
VEMUS	http://www.tehne.ro/projects/vemus_virtual_music_school.html	R	4.96	2.60
iMaestro	http://www.i-maestro.org/	R	4.19	2.82

Table 3. Overview of aggregated scores for each platform.

Platform	depth	object tags-cats	depend. objects	difficulty organization	content organization	fixed guide-through	sequence strategy	assessment?	multiple-choice	student model	goals
Coursera	3	0	0	0	1	1	0	1	1	0	0
Playground sessions	2	0	0	1	1	1	0	1	0	0	0
Berklee Pulse	2	0	1	1	1	1	0	1	0	0	0
Musictheory.net	3	0	0	0	1	1	0	1	0	0	0
Allversity Guitar Beginners	3	0	0	0	1	1	0	0	0	0	0
PRAISE	1	1	0	0	1	1	1	0	0	1	1
VEMUS	2	1	0	0	1	0	0	1	1	1	0
iMaestro	2	0	0	0	1	0	0	1	1	1	0
StackExchange	1	1	1	0	0	0	0	0	0	0	1
YouTube	3	0	1	0	0	0	0	0	0	0	0
iReal Pro	1	0	0	0	0	0	0	0	0	0	0
Theta Music Trainer	0	0	0	1	0	0	0	1	0	0	0
Udemy "Learn free music theory"	2	0	0	0	1	1	0	0	0	0	0
Howtoplaypiano.ca	1	0	0	0	1	1	0	1	0	0	0
Music 101: Intro to Music	2	0	0	0	1	1	0	1	0	0	0
BBC Sing	1	0	0	0	1	1	0	0	0	0	0
Yale "Music 112: Listening to music"	3	0	0	0	1	1	0	0	0	0	0
MIT Introduction to Music Composition	3	0	0	0	1	1	0	0	0	0	0
Online music theory tutor	1	0	0	0	1	1	0	1	0	0	0
MIROR	1	0	0	0	0	0	0	0	0	1	1
Piano Suite Premier	2	0	0	0	1	1	0	1	0	0	0
Piano Marvel	2	0	1	0	1	1	0	1	0	0	0
Artist Works Music & Arts campus	3	0	0	1	1	1	1	1	0	1	0
PianoStreet	2	0	0	0	0	0	0	0	0	0	0

Table 4. Overview of indicator scores for curriculum design for each platform.

Platform	Domain model	Feedback insight	Qualitat. feedb	Gener. exerc	Game / play	Info on time	Adapted skill level	Diag & repairs (exercises)
Coursera	0	0	1	0	0	1	0	0
Playground sessions	0	0	0	0	1	1	0	0
Berklee Pulse	0	0	0	1	1	1	0	0
Musictheory.net	0	0	0	1	1	1	1	0
Allversity Guitar Beginners	0	0	0	0	0	0	0	0
PRAISE	1	1	1	0	0	1	0	0
VEMUS	1	1	0	1	0	1	1	0
iMaestro	0	1	1	1	1	1	0	0
StackExchange	0	1	1	0	0	1	1	0
YouTube	0	0	0	0	0	0	0	0
iReal Pro	0	0	0	0	0	0	0	0
Theta Music Trainer	0	0	0	0	1	1	1	0
Udemy "Learn free music theory"	0	0	0	0	0	0	0	0
Howtoplaypiano.ca	0	0	0	0	0	0	0	0
Music 101: Intro to Music	0	0	0	0	0	0	0	0
BBC Sing	0	0	0	0	0	0	0	0
Yale "Music 112"	0	0	0	0	0	0	0	0
MIT Music Composition	0	0	0	0	0	0	0	0
Online music theory tutor	0	0	0	0	1	1	1	0
MIROR	0	0	1	1	1	1	1	0
Piano Suite Premier	0	0	0	1	1	1	0	0
Piano Marvel	0	0	0	0	1	1	0	0
Artist Works Music	0	1	1	0	0	0	0	1
PianoStreet	0	1	1	0	0	1	1	0

Table 5. Overview of indicator scores for personalized learning.

References

[1] I Elaine Allen and Jeff Seaman. *Changing Course: Ten Years of Tracking Online Education in the United States.* ERIC, 2013.

[2] Allen, I. Elaine and Seaman, Jeff. Grade Change. Technical report, Babson Survey Research Group and Pearson and Sloan, January 2014.

[3] John Baer. Grouping and achievement in cooperative learning. *College Teaching*, 51(4):169–175, 2003.

[4] William Bauer. Technological pedagogical and content knowledge for music teachers. In *Society for Information Technology & Teacher Education International Conference*, volume 2010, pages 3977–3980, 2010.

[5] William I Bauer. *Music learning today: Digital pedagogy for creating, performing, and responding to music.* Oxford University Press, USA, 2014.

[6] George A Beauchamp. Curriculum theory. *ERIC*, 1968.

[7] Katrien Beuls. *Towards intelligent MOOCs*, chapter 11. IOS press, 2015.

[8] Judith Bowman. *Online Learning in Music: Foundations, Frameworks, and Practices.* Oxford University Press, 2014.

[9] Brown, Alexandria. Moocdemic starts spreading September 29th. http://www.examiner.com/article/moocdemic-starts-spreading-september-29th, September 2014.

[10] Button, Keith. U of OK for-credit courses offered as MOOCs. http://www.educationdive.com/news/u-of-ok-for-credit-courses-offered-as-moocs/334229/, November 2014.

[11] Evi Chryssafidou and Anders Askenfelt. VEMUS Project - Evaluation Report v3, Deliverable D2.3v3. Technical report, VEMUS Consortium, IST-27952, 2009.

[12] Gráinne Conole and Panagiota Alevizou. A literature review of the use of web 2.0 tools in higher education. *A report commissioned by the Higher Education Academy*, 2010.

[13] M. Csikszentmihalyi. *Beyond boredom and anxiety: experiencing flow in work and play.* Cambridge University Press, 1978.

[14] Richard J Dammers. Utilizing internet-based video-conferencing for instrumental music lessons. *Update: Applications of Research in Music Education*, 2009.

[15] Pierre Dillenbourg et al. Collaborative-learning: Cognitive and computational approaches. Technical report, Elsevier, 1999.

[16] Fisher, Aaron and Anderson, G. Brooke and Peng, Roge and Leek?, Jeff. A randomized trial in a massive online open course shows people don't know what a statistically significant relationship looks like, but they can learn. *PeerJ*, October 2014.

[17] D Randy Garrison. *E-learning in the 21st century: A framework for research and practice.* Taylor & Francis, 2011.

[18] "Online Master of Science in Computer Science". http://www.omscs.gatech.edu/, 2014.

[19] Judi Harris and Mark Hofer. Instructional planning activity types as vehicles for curriculum-based tpack development. In *Society for Information Technology & Teacher Education International Conference*, volume 2009, pages 4087–4095, 2009.

[20] "Online Evolution". http://harvardmagazine.com/2014/09/online-evolution, September 2014.

[21] Daniel CA Hillman, Deborah J Willis, and Charlotte N Gunawardena. Learner-interface interaction in distance education: An extension of contemporary models and strategies for practitioners. *American Journal of Distance Education*, 8(2):30–42, 1994.

[22] IBM 2010 Global CEO Study. http://www-03.ibm.com/press/us/en/pressrelease/31670.wss, May 2010.

[23] iMaestro website. http://www.i-maestro.org/.

[24] Jhee, Catherine. Digital Games in the Classroom: A National Survey. http://www.joanganzcooneycenter.org/2014/06/09/digital-games-in-the-classroom-a-national-survey/, June 2014.

[25] John-Steiner, V. *Creative Collaboration.* Oxford twenty-first century approaches to literature. Oxford University Press, 2006.

[26] Mauritz Johnson. Definitions and models in curriculum theory. *Educational Theory*, 17(2):127–140, 1967.

[27] Kahn, Bob. Genetics in Minecraft: Wool, Mendel and Brownies. http://www.middleschoolminecraft.com/2013/04/08/genetics-in-minecraft-wool-mendel-and-brownies/, April 2013.

[28] Kim, Joshua. Here Come the Data Scientists. https://www.insidehighered.com/blogs/technology-and-learning/here-come-data-scientists, August 2014.

[29] Knox, Jeremy and Bayne, Sian and MacLeod, Hamish and Ross, Jen and Sinclair, Christine. MOOC pedagogy: the challenges of developing for Coursera. https://newsletter.alt.ac.uk/2012/08/mooc-pedagogy-the-challenges-of-developing-for-coursera/, August 2012.

[30] Kyung Hee Kim. The Creativity Crisis: The Decrease in Creative Thinking Scores on the Torrance Tests of Creative Thinking. *Creativity Research Journal*, 23:4:285–295, 2011.

[31] Laurillard, Diana. Five myths about MOOCs. http://www.timeshighereducation.co.uk/comment/opinion/five-myths-about-moocs/2010480.article, January 2014.

[32] Lynch, Michael. CMU Gets $5 Million Grant To Study Learning. http://wesa.fm/post/cmu-gets-5-million-grant-study-learning, October 2014.

[33] Terry Mayes and Sara De Freitas. Review of e-learning theories, frameworks and models. *JISC e-learning models desk study*, 1, 2004.

[34] Barbara Means, Yukie Toyama, Robert Murphy, Marianne Bakia, and Karla Jones. Evaluation of evidence-based practices in online learning: A meta-analysis and review of online learning studies. *US Department of Education*, 2009.

[35] Musical interaction relying on reflexion - miror. http://www.mirorproject.eu/.

[36] Punya Mishra and Matthew Koehler. Tecnological pedagogical content knowledge: A framework for teacher knowledge. *The Teachers College Record*, 108(6):1017–1054, 2006.

[37] Moe, Rolin. The MOOC Problem. http://www.hybridpedagogy.com/journal/mooc-problem/, May 2014.

[38] Michael G Moore. Learner autonomy: The second dimension of independent learning. *Convergence*, 5(2):76–88, 1972.

[39] Nichols, Valerie. Should you use MOOCs in learning? http://www.trainingzone.co.uk/feature/od/should-you-use-moocs-learning/187882, November 2014.

[40] Obama, Barack. http://www.whitehouse.gov/photos-and-video/video/2011/03/09/president-obama-education-techboston#transcript.
[41] Kenneth H Phillips. Graduate music education. *Research and Issues in Music Education*, 6(1), 2008.
[42] F. J. Rees. Distance learning and collaboration in music education. In Richard Colwell and Carol Richardson, editors, *The new handbook of research on music teaching and learning: A project of the Music Educators National Conference*. Oxford University Press, 2002.
[43] Romero, Cristóbal and Ventura, Sebastián. Educational data mining: a review of the state of the art. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 40(6):601–618, 2010.
[44] Heikki Ruismäki and Antti Juvonen. The new horizons for music technology in music education. *The Changing Face of Music Education. Music and Environment*, pages 98–104, 2009.
[45] Sadler, Philip M and Good, Eddie. The impact of self-and peer-grading on student learning. *Educational assessment*, 11(1):1–31, 2006.
[46] Miikka Salavuo. Social media as an opportunity for pedagogical change in music education. *Journal of Music, Technology & Education*, 1(2-3):121–136, 2008.
[47] Luc Steels, Mark d'Inverno, Francois Pachet, and Carles Sierra. Praise - practice and performance analysis inspiring social education, 2013.
[48] M. Tokoro and L. Steels, editors. *The Future of Learning: Issues and Prospects*. IOS Press, 2003.
[49] Cheryl L Ward and Susan N Kushner Benson. Developing new schemas for online teaching and learning: Tpack. *MERLOT Journal of Online Learning and Teaching*, 6(2):482–490, 2010.
[50] Waxman, Olivia B. MinecraftEdu Teaches Students Through Virtual World-Building. http://techland.time.com/2012/09/21/minecraftedu-teaches-students-through-virtual-world-building/, September 2012.
[51] Webley, Kayla. MOOC Brigade: Can Online Courses Keep Students from Cheating? *Time*, November 2012.
[52] Peter R Webster. Computer-based technology and music teaching and learning: 2000–2005. In *International handbook of research in arts education*, pages 1311–1330. Springer, 2007.
[53] Etienne Wenger. *Communities of practice: Learning, meaning, and identity*. Cambridge University Press, 1998.
[54] Westervelt, Eric. The Online Education Revolution Drifts Off Course. http://www.npr.org/2013/12/31/258420151/the-online-education-revolution-drifts-off-course., December 2013.
[55] "Distance education". http://en.wikipedia.org/wiki/Distance_education.
[56] "Educational Data Mining". http://en.wikipedia.org/wiki/Educational_data_mining.
[57] Beverly Park Woolf. *Building intelligent interactive tutors: student-centered strategies for revolutionizing e-learning*. Morgan Kaufmann Publishers, 2009.
[58] Fareed Zakaria. When big data meets education. http://globalpublicsquare.blogs.cnn.com/2014/11/15/when-big-data-meets-education/, 11 2014.

Part II

The Role of Feedback

This page intentionally left blank

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-41

Chapter 3.
Social Feedback as a Creative Process

Mark d'Inverno and Arthur Still
Department of Computing, Goldsmiths, University of London,
dinverno@gold.ac.uk, awstill@btinternet.com

Abstract. Arguably one of the most important activities of a university is to provide environments where students develop the wide variety of social and intellectual skills necessary for giving and receiving feedback. We are not talking here about the kinds of activity typically associated with the term "feedback" — such as that which occurs through individual course evaluation questionnaires or more universal systems such as the National Student Survey, but the profoundly creative and human act of giving and receiving feedback in order to validate, challenge and inspire. So as to emphasise we are talking about this kind of feedback, we coin the term "creative feedback" to distinguish it from the pre-conceived rather dreary compliance-inflected notions of feedback and set out in this paper to characterise its qualities. In order to ground and motivate our definition and use of "creative feedback" we take a historical look at the two concepts of creativity/creative and feedback. Our intention is to use this rich history to motivate both the choice of these two words, and the reason to bring them together. In doing so we wish to emphasise the characteristics of an educational philosophy underpinned by social interaction. By describing those qualities necessary to characterise creative feedback this paper sets out an educational philosophy for how schools, communities and universities could develop their learning environments. What we present here serves not only as a manifesto for designing learning environments generally, but as a driver for designing technologies to support online social learning, as captured in the concept of social Moocs [70]. Technology not only provides us with new opportunities to support such learning but also to investigate and evidence the way in which we learn and the most effective learning environments.

Keywords. Feedback, creative, creativity, learning, technology

1. Introduction

When the word feedback is mentioned in universities — as happens now with increasing frequency — there are usually one or two winces around the room. The problem it is a word that has become associated with compliance, with checking competency, with measurement and judgement, with having to go through the motions of various government or funding body processes and, perhaps too, with feeling beholden to open up channels of communication so as to hear things that we would rather not have to hear. This is a pity, and especially so at universities, because feedback is central to learning. Not just to learn a discipline, but to learn about the way we are, to learn about the way we think, to learn about the way we interact and about the way in which we produce and value our work. Whether that work is an analytical or interpretive essay, whether it is a poem or a composition, whether it is a new performance or a new artwork, it is only through actively seeking feedback both from others and from ourselves that we learn.

At one level it is clear that without the on-going feedback that we sense and perceive from our environment we could not operate or survive. Without basic perceptual acts such as seeing, hearing and touching we couldn't function for very long. However, feedback is also necessary to experience ourselves as social beings, and especially to understand and investigate the process of social interaction between individuals. Sometimes the communication from one human to another is like an experiment whose result is evidenced by the feedback perceived from the other [1]. For example, shouting "hello?" to check whether anyone is at home, the result might be the perception of a response like "I'm in the kitchen!" or complete silence. This is an example of a simple feedback loop at work providing evidence for a model of the world. At the other extreme feedback loops can be continuous and extremely complex, and often below conscious awareness such as when two jazz musicians are improvising together [2]. In all cases feedback is the way in which we understand the world we are in, and learn about our physical and social place within it.

Suppose you are learning to play music, for example. If you play a piece of music then the only way you can know how it was heard and experienced by others is to get their feedback on your performance. This feedback will be absolutely critical if you want to understand how you can improve yourself as a performer. Of course in any performance sustained self-feedback is critical too and musicians are skilled enough to give themselves this on-going and continuous feedback as they play. In addition to this, musicians have the option of recording performances and listening to them later in order to provide an entirely new perspective. The distance created in time and space, and moving from performer to listener, provides new opportunities for fresh insights on how to improve ones own performance. In addition, through an understanding of how we come across to others, we can often best advance the quality and precision of the feedback we give ourselves.

If we accept the need for building communities of feedback the issue then becomes how to build the right kinds of learning environments. If students can develop their own skills in giving and receiving feedback at school and university, then they will gain confidence in giving and receiving feedback from friends, colleagues, press and audiences too. Education environments should enable an exploration of how peers and tutors perceive essays, performances, software and artworks and in turn, how we all learn to be open to the feedback from others.

This philosophy is very strong in the Art department at Goldsmiths, where the emphasis is very much focused on developing communities of feedback. This department is especially interesting because of its reputation for producing world-class artists that have become important cultural and creative pioneers in the UK.[1]

In our observations, first, second and third year undergraduates come together weekly in order to give feedback on a small selection of undergraduates work. The students clearly worked as a group in balancing praise and criticism, combining the emo-

[1] (Damien Hirst, Malcolm McClaren, Mary Quant, Lucien Freud and Anthony Gormley are all alumni of the Art department. Other alumni include Laurie Provoust who currently holds the Turner prize and Steve McQueen who won a Bafta and Oscar for best film with "12 years a slave". The question to us is whether developing communities of creative feedback is the key to the Art department's success.)

tional and analytical, and moving from the sociological to the political. In all these open conversations students are learning about how to give and receive feedback to each other and understanding the ever present gap between any intention behind an artwork, and the perception by others. One of the most fascinating aspects observed in these sessions was the ability of students to take a sufficient emotional distance in order to be open to feedback, and to experience it freely without personalising anything. This ability is not only key in terms of learning how others experience their work but becomes an important skill for artists moving into a professional sphere with the free-for-all comment and criticism that social media now encourages.

Arguably then, a learning institution's key objective is to provide the kind of supportive and trusting environments where students can develop their ability to give and receive feedback in a culturally-aware, sensitive, mindful, critical and challenging way. We certainly think so, and would like a label to describe the kind of feedback we have in mind, and for this we choose the term "creative feedback". In this paper we provide a historical account of the notions of creative and creativity in order to justify the use of this term in an educational context. Moreover, by using this term explicitly the hope is we can rescue the concept of feedback from its often rather dreary compliance-inflected interpretation.

In what follows we will call upon our experience as educators spanning mathematics, psychology, psychotherapy, music and computer science, to try to explain what we mean by creative feedback and to justify our use of this term. To do this we need to take a brief historical look at the concepts of "creative" (and the related "creativity") and "feedback" — particularly, though not exclusively, in an education context — in order to explain exactly what we mean by these terms and why we are bringing them together specifically. The aim of the historical analysis is to give currency to the use of the term and the underlying manifesto for learning. We clearly need to be mindful of using the word "creative" when it is used so loosely, and for so many different educational, marketing and political reasons. We not only have creative writing and creative learning but now we have creative musicianship, creative computing and creative financing, not to mention the growing importance given to "creative industries" and economic arguments about why they are such an important part of our future. The word is in danger of being no more than what is approved of, and we wish to recover an older and fuller meaning for our purposes.

1.1. Aims

In this paper we set out to characterise creative feedback as the basis of an educational philosophy that is inspired by the American psychologist, philosopher, and educationalist John Dewey. The idea that follows naturally from this is that we structure schools, learning groups and universities as "communities of discovery". There are a number of motivating factors for the work in this paper described next.

The first is the desire to build educational environments (which include online environments) that give more people access to developing "creative feedback" skills. Creative feedback belongs to what Dewey called "creative intelligence" which is a part of all human thinking and is available to everyone. A strong part of our individual learning journey is gaining an understanding how others see us. The way we think, the way we behave, what we produce. This understanding is such a crucial part of learning that we want to build environments that encourage students to be aware of how others see them. As George Herbert Mead wrote, *"the individual mind can exist only in relation to other minds with shared meanings"* [3, p5]. If this is true, the relation to other people is grounded within a framework of feedback and the individual mind can only exist within such a framework.

Next, we want to emphasise that "creativity" depends on feedback from the world rather than being something that is an intrinsic quality that resides within individuals. It depends on feedback both in the act of creation itself, and also the social feedback that is received once it is made available to others (which may or may not amount to acclamation as great art).

As stated above feedback is not often seen as a creative endeavour but rather as being quite mechanical (tick boxes and scores) and about compliance (such as is often the case when making module feedback forms available to students). The impact of this notion of feedback on tutor/tutee relationships can often be dire. We explicitly introduce "creative feedback" to mitigate against this commonly held view of feedback and, in addition, to move away from another commonly held conception about feedback that it only exists in terms of praise and punishment. Furthermore, we want to emphasise how we are immersed in feedback as biological and social beings and we wish any definition to encompass this.

Most educationalists like us want to promote effective education as available to everyone rather than a middle-class luxury and technology clearly has an important role here. However, technology also provides opportunity to bring communities of learners together and, moreover, serve as a test-bed from which we can start to evidence the benefits of social learning over the individual, rote-learning and exam-based methodology that so dominates current political thinking. It also provides us with exciting new possibilities for understanding the way in which we learn. One of the drivers in our own research, for example, is to develop learning analytics and methodologies that can enable us to correlate creative feedback with learning.

The ability to use technology to understand and support social learning depends on whether we can construct systems that encourage humans to give and receive creative feedback. In order to achieve this we need participatory design methods working with a variety of user groups in order to design software that can support creative feedback across a whole range of disciplines (e.g. poetry, music, design, digital art). We believe a historical and educational underpinning is necessary to drive the principled design of such systems that not only support creative feedback but also allow mixed human and computational societies. One of the practical questions that we are addressing in the design of novel education systems that enable social learning is how to build autonomous artificial systems that can help exemplify creative feedback in a learning community.

2. A History of Creativity and Feedback

2.1. The Education Wars

Ever since people started arguing about education, there has been an angry debate that is still not resolved, and is especially marked today in England. On the one hand, the Secretary of State for Education, Michael Gove, crusades for even more frequent and stringent examinations and inspections in the State-based schools, creating what his critics call *"exam factories"* [4], designed to compete with the dauntingly efficient exam factories of the Far East. And on the other hand the popular educationalist Sir Ken Robinson speaks for many when he condemns such an approach for undermining creativity, which is, according to this view, the true goal of democratic education. It may be hard to define creativity, but everyone agrees that it is a good thing, and that it is not fostered by an exclusive focus on training students for success in exams. The emphasis on exam factories may even be self-defeating, since there are studies showing that the success of children in China and Japan depends more on the early nurturance of sociality, than on forced study and rigorous examinations [5]. That is, working more like what Coffield called *"communities of discovery"* than *"exam factories"*, so perhaps Gove is taking us *"ever faster down the wrong road"* [6].

2.2. Background to the Conflict

This quarrel occurs at every level of education, from toddlers to adults, and it reflects different views on the nature of children. At one extreme is the active child, full of wonder and curiosity at the world, who needs only skilled guidance from the teacher to flower into a civilized and creative adult. At the other is the resistant child, lazy and easily distracted, whose motivation and attentiveness require firm moulding and sometimes medication in order to learn lessons and become a good citizen. Around 1900 these extremes were given psychological and educational form by two prominent American thinkers [7], and this set the scene for many of the debates on education during the coming century. In the active, curious child camp sat the philosopher, educationalist and psychologist, John Dewey, the great champion of American pragmatism, which is a philosophy based on doing rather than thinking; in the other camp sat Edward Thorndike, famous throughout the 20^{th} century for his puzzle box experiments with cats published in 1898 [8] in which he claimed to show that cats are incapable of reason and learn only through trial and error. During the second half of the 20^{th} century both camps contributed to the new interest in creativity, which has now become a massive and well-funded research industry in Europe especially in relation to technology.

In this paper we aim to show how technology can contribute to the fostering of creativity in education in a way that can satisfy both the jeremiads of Professor Robinson and the ministerial anxieties of Michael Gove. But first we need to be clear about what kind of learner we have in mind, Dewey's or Thorndike's, since this determines what we mean by creative and creativity, and the deployment of these terms has provided a map of the hidden agendas of Psychology and Educational Theory during the 20^{th} century.

2.3. E. L. Thorndike: Connectionism, Stimulus-Response and the Importance of Measurement

In 1911 Thorndike published his puzzle box experiments in *Animal Intelligence*, and developed the theory that learning is initially guided by random trial and error learning, rather than rational intelligence. For Thorndike and later many Behaviourists, the unit of behaviour was the stimulus response (S-R) connection, treated as a kind of reflex. Thorndike's view was that learning takes place by establishing connections in the brain and these connections are stamped in through a system of reward and punishment.

Applied to education it was argued that the randomness of the trials in initial learning showed that little is to be gained by relying on the prior capacities of the novice learner.

Connections were treated as "atoms of the mind", and Thorndike speculated that *"the vague gross feelings of the animal sort might turn into the well-defined particular ideas of the human sort, by the aid of a multitude of delicate associations"* [9, p289]. This is Thorndike's Connectionism, and it has been one of the main models guiding studies of learning throughout the 20th century, though it was quickly found that the S-R scheme needed to be extended to S-O-R [10]. In this extended scheme O refers to the state of the organism, which is made up of many variables or factors, including prior knowledge (the multitude of delicate associations), motivation, attentiveness, intelligence and many other variables.

During the second half of the 20th century computers became the new model of the mind, and the language for describing "a multitude of delicate associations" became increasingly sophisticated, eventually leading to a new brand of Connectionism as a model for perception and learning [11]. But even in its most sophisticated form, it is still about the selection of successful acts and the "stamping out" of "profitless" [9, p283] acts by reward and punishment. Nowadays we speak of input and output of information rather than S-R, but whatever the cognitive complexity of what goes on in between, a basic linear structure remains, with the environment operating on the organism, rather than the organism on the environment.

But Thorndike was not only one of the founders of S-R theory, he was also a pioneer of mental testing as a way of classifying individuals for social control, and therefore for assigning numbers to the "O" variables in the S-O-R scheme. Thorndike greatly admired the work of Darwin's cousin Francis Galton (1822-1911) who spent much of his life studying and measuring human variation and its genetic basis after reading *Origin of Species*. As part of this interest Galton became the first to use questionnaires and statistics for the measurement of human differences and Thorndike in turn became a champion of measurement in Psychology and Education. In 1904 he published *An Introduction to the Theory of Mental and Social Measurements* [12] which introduced students to the new statistical methods that were to dominate the scientific practice of Psychology.

2.4. Deweyan Inquiry

The contrasting philosophy was that of John Dewey, who was one of the first to acknowledge the value of Galton's statistical discoveries [13] but had little faith in the value of measuring the worth of individual human beings [14]. He believed effective education is powered by the child's spontaneous curiosity about the world and is social, taking place in "a community held together by participation in common activities" [15,61]. This social setting generates "inquiry", a process as natural as breathing in all animals. Inquiry is an ongoing process that reveals "novelty", which in turn becomes the spur to further inquiry.

In 1896 Dewey had made the revolutionary step of taking the basic S-R reflex studied in the laboratory by physiologists, not as the simple arc of Thorndike, but as a circular structure with neither stimulus nor response being dominant over the other. He argued that the S-R reflex is not an isolable molecule of behaviour, but is inseparable from an ongoing process involving what 50 years later would be called feedback.[2] Dewey was not a laboratory psychologist, and unlike Thorndike's S-R, his scheme did not lend itself to precise control, since it required freedom of action for optimal learning to take place.

The main concern for the teacher therefore is to guide this action toward educational goals, and to avoid stifling freedom through the indiscriminate "stamping out" of what Thorndike referred to as "profitless" acts. For Dewey these "profitless" acts are part of what he called inquiry and to stamp them out is to suppress inquiry and to stunt human development.

2.5. Who Has Won?

In Psychology and in Education, Thorndike has won hands down: *One cannot understand the history of education in the United States during the twentieth century unless one realises that Edward L. Thorndike won and John Dewey lost* [16, p185].

But as Lagemann goes on to point out, Dewey paradoxically remains a significant figure in education, dominating discussion in schools of education, and pointing to an ideal, even if it is Thorndike who prevails in practice. But occasionally an indirect Deweyan light shines through. A possible example of this was the dramatic reception in the West of Vygotsky's Zone of Proximal Development (ZPD). Dewey had a strong influence on Russian education in the 1920's when Vygotsky was developing his ideas [17]. Vygotsky had certainly read Dewey's work [18, p53], and there is a close affinity with Dewey's ideal of "a community held together by participation in common activities" [15, p55]. ZPD contrasted the child's developmental level when measured by conventional tests, with the level shown under adult or peer guidance [18, p86] where the ability to follow and imitate comes into play: *"using imitation, children are capable of doing much more in collective activity or under the guidance of adults"* [19, p88]. This presupposes *"a specific social nature and a process by which children grow into the intellectual life of those around them"* [18, p88], which comes close to the collective learning through inquiry described by Dewey. In 1966 Bruner [20] introduced the word "scaffolding" to describe what is going on in ZPD, but this has been often been limited to the capacity to benefit from adult help [21], rather than from the more general sociality of "collective activity", which leads to a form of "social constructivism" [22]. Like an education based on Deweyan inquiry, ZPD in our interpretation goes very deep, and its effects, unlike those of scaffolding (if we take the metaphor literally), cannot be removed once the construction is complete.

In Psychology too, Dewey has been lurking in the background, and his influence became more apparent once the notion of feedback spread after the publication of Norbert Wiener's *Cybernetics* [23]. Later, in 1960, *Plans and the Structure of Behavior* [24] appeared, and brought together feedback of information (rather than reward and punishment) with some of the early influences on Artificial Intelligence. These included Chomsky's generative grammar [25] and Newell, Shaw and Simon on problem solving in com-

---

[2]Thorndike's S-R connectionism also involved a rudimentary form of feedback. Reward and punishment applied to isolated S-R connections are feedback. But Dewey seemed to have in mind what we now think of as a self-organising system, in which the parts, which we may for convenience label stimulus, response, feedback, etc., cannot usefully be isolated and studied as "laboratory preparations" outside the system. The knowledge gained by an inquiring child involves, not a changing array of S-R connections, but an evolving place within a system that includes its social and physical environment.

puters [66]. The result was the TOTE (test operate, test exit), introduced as a unit of behaviour to replace the S-R model, and the authors were quick to recognise that this was similar to what Dewey had proposed in his 1896 reflex arc paper [26,27, p30].

More generally, affinity with the Dewey scheme rather than Thorndike's shows itself when the organism, animal or human, is treated as essentially in the world, active and subject to continuous feedback as it acts, rather than a static processor of information. Examples of this Deweyan scheme are Gibson's sensori-motor systems as a model for perception [28]; the move in Robotology from cognitive representions to a focus on sensori-motor activity [29]; Jean Lave's Situated Learning [30]; and more recent work in Psychology and Philosophy on Situated Cognition [31].

2.6. Formative Assessment and Feedback

In one respect - through the notion of formative assessment - the Deweyan influence penetrated deep into the heartlands of Thorndikean territory, measurement and educational testing.

The psychologist L.L.Thurstone studied at Chicago with a close colleague of Dewey's, George Henry Mead, and spent most of his career there. Early on in his career he proposed a Deweyan model of ongoing behaviour as an alternative to the S-R scheme [32]. But his main achievements were in test theory and a more careful analysis than was usual of what is typically meant by measurement in Psychology [33]. Lee Cronbach, whose PhD was also from Chicago, continued this critical tradition within psychological measurement. His work with Meehl on Construct Validity [34] showed the limitations of psychological testing, since it measures constructs rather than reality. And he recommended that assessment be part of the learning process, rather than a test given after the learning is over [35]. Later this was labelled "formative" by contrast with the conventional "summative" assessment [36], the one made by tests after the course has ended; whereas formative assessment is the one made during the course, designed as part of the learning process. It is closer therefore to a Deweyan rather than a Thorndikian philosophy of education, and the formative assessor joins "a community held together by participation in common activities" [15, p55]. Formative assessment involves what came to be called formative feedback. In formative feedback the student is given ongoing information about performance, and the term has replaced the concepts of reward, punishment and reinforcement. But the old S-R scheme dies hard, and many of the experiments reported on formative feedback seem quite similar to those by Thorndike and others of 80 years ago [37]. They are a long way from the feedback of a sensori-motor system that is the necessary vehicle for Deweyan inquiry. This same pattern — an apparent massive victory by the Thorndike camp, yet a persistent critical or subversive presence from the Deweyans — exists in the field of creativity, where the difference between the two viewpoints is especially marked and important given that the concept of creativity is so dominant in educational discourse.

2.7. Creative Intelligence

In literature on Creativity, which spans many disciplines and is now remarkably large and increasing every year, two distinct points of view about its nature have remained unchanged. The first is that it is a perfectly ordinary and basic property of all human and perhaps even animal behaviour. The second is that it is a puzzling and wonderful property of the human mind that has given rise to all great human achievements.[3] The reason for this strange contradiction between the two meanings, which seems to have gone largely unnoticed, may be because the modern word "creativity" derives from two distinct ways of thinking about novelty and innovation in the world. The first of these, which sees creativity as the basic process of every mind, belongs to the Deweyan view. The second, which came later, sees creativity as a marvellous addition to the mechanical processes of ordinary thinking; this belongs to the Thorndikean view.

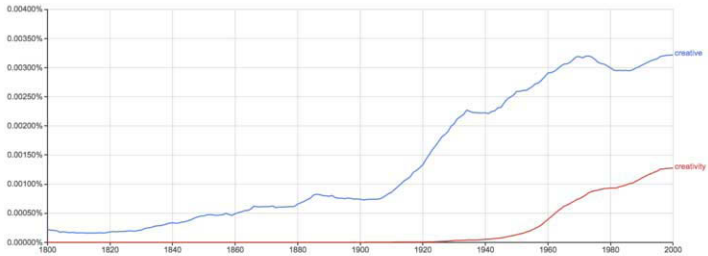

Figure 1. Creative and Creativity in Google's nGram

As the diagram above suggests, the popularity of words like "creative" and "creativity" is only quite recent. Originally both words were the prerogative of God, who was unique in being able to make something (the world) out of nothing. This is what creation meant, making something out of nothing. With this in mind, "Creative" (though not creativity) was occasionally extended to women giving birth and in the 19th century to refer to the divine and mysterious work of poets and artists.[4] This can be seen clearly in the diagram above.

But after the widespread acceptance of the Theory of Evolution by the end of the 19th century, the world itself could be seen as creative through variation and selection, with no help from God. This is how it is used in the title of Bergson's *Creative Evolution* [41] which was first published in French in 1907, and then translated into English four years later[5]. This was a book that was widely discussed, especially in the pragmatist circles around William James in Harvard and John Dewey in Chicago.

Dewey's *Creative Intelligence* was published later in 1917, and the word "creative" in the title was not being used to pick out one kind of intelligence amongst others, but to emphasise that human intelligence is inherently creative through a natural process of deliberate variation and invention. This could be the herald of a new beginning for education, since according to the traditional philosophies, *"If ever there was creation it all took place at a remote period. Since then the world has only recited lessons."* [42, p23]. Dewey thought that reciting lessons is a way of suppressing the variation that is necessary for creative intelligence to flourish. There was nothing divine about Dewey's view of creative thought, and he made little use of the popular concept of genius, instead

3 "Creativity is consensually viewed as one of the most remarkable characteristics of the human mind." [40, p147]. Creativity "is the humble human counterpart of God's creation" [38, p4].

4 "But this I know; the writer who possesses the creative gift owns something of which he is not always master—'something that at times strangely wills and works for itself'." Charlotte Bront in editorial preface to 1850 edition of *Wuthering Heights* [39, p 1iii].

5 Translation of Bergson's *L'Évolution créatrice* from 1907 as *Creative Evolution* in 1911 [4]

seeing art and creativity as present in the most mundane activities: *"The sources of art in human experience will be learned by him who sees how the tense grace of the ball-player infects the onlooking crowd; who notes the delight of the housewife in tending her plants, and the intent interest of her goodman in tending the patch of green in front of the house"* [43, p3].

In this philosophy, education involves social control, but not via rules dictated by authority. Instead Dewey took as a benign paradigm of social control that of children playing games, in which the control is not from on high, but is naturally social from "a community held together by participation in common activities" [44, p55]. This underlies his practical experiments in education in the experimental schools he set up first in Chicago, later at Columbia University.

2.8. Creativity

The modern word "Creativity" came into play a little later than "creative," in the mid 1920's [44]. In 1924, around seven years after Dewey's Creative Intelligence was published, the mathematician and philosopher Alfred North Whitehead was invited to Harvard, where he developed the process philosophy for which he is best known. At the centre of this philosophy was his concept of creativity, a term he coined from the Medieval Latin "creare". [18, p208]. This was his word for the evolution of forms or species. Darwin had shown how this could be a property of organic evolution, and Whitehead applied the same basic structure (variation, and a means of fixing change) to the universe as a whole. It was his metaphysical principle through which entities are created out of flow (*"all things flow"* [45, p208]) which is more basic than the things that we experience. New forms (the solar system, new species) emerge and creativity is the power that enables this to happen. Dewey read this as a universal generalisation of his own views of human invention, managed by creative intelligence out of variation, and wrote approvingly about Whitehead and his ideas of creativity in 1937 [46]. On this view, there is nothing special about creativity. It is a basic principle of the world, and human creativity is no more than a reflection of this.

2.9. From Creativity to Social Creativity

Dewey's friend and colleague the social psychologist G.H. Mead had contributed one of the chapters in Dewey's Creative Intelligence of 1917 writing, *"The individual in his experiences is continuously creating a world which becomes real through his discovery"*. [47, p210] After reading Whitehead, he used the word "creativity" in his lectures during the 1920's, [47, p325], and it appeared in his best known book *Mind, Self and Society* [48] which was widely read.

There Mead described how any individual self is constituted by the social and physical environment it inhabits, but at the same time affects the environment in which the it is situated. More generally, the organism is partly determined by its environment, but also *"is determinative of its environment"*, a more general version of the circular process described by Dewey [49]. Thus the word "creativity" is will have been familiar to the many readers of Mead and Dewey, and they would have had a common understanding that there was nothing special about it, not linked to genius but essential for the thinking of every human being and animal.[6] [65, p10-11]

2.10. Creativity as Faculty

But when creativity re-emerged in 1950 [50] it had a different meaning, and came from a different tradition of Psychology, that of Psychological measurement, therefore closer to Thorndike than to Dewey. It was not about creativity as the generation of change and novelty in the world, but referred instead to a personality characteristic. Launched by J.P. Guilford in 1950 in a presidential address to the American Psychological Association, he started by expressing astonishment at the lack of work on Creativity. He made no mention of Whitehead, Dewey or Mead, and based his concept of creativity on Factor Analysis, discovered by Charles Spearman [51]. Spearman had actually written a book called Creative Mind in 1930 [52], in which the word "creativity" appears, but it is not referred to by Guilford though he is likely to have known it. Spearman was a colleague of Whitehead's at UCL for several years before Whitehead left for Harvard, and may have picked the word up from him.

By partitioning similar correlations in tables from a large number of tests, Spearman had shown how to extract distinct factors of the mind, like intelligence, perseverance, memory and so on, and now creativity, which can be used to form part of the O in the S-O-R scheme. By 1950 Factor Analysis had reached a high level of sophistication, and Guilford had isolated a factor he called Creativity, based on his test of Convergent and Divergent thinking. Convergent thinking is conventional problem solving, converging on the correct solution, divergent is open ended and was thought to allow the free play of imagination, with questions like "in what different ways can you make use of a brick" Later many other tests of creativity were devised including Torrance's Incomplete Figure Test [53] tests of insight, similar to Duncker's classic candle problem [54] and of "remote associations" by Mednick et al [55].

2.11. The Creativity Bandwagon

The vastness of the bandwagon launched by Guilford has been extraordinary, and cannot be explained only by the happy Utopian vision offered by the definition that runs throughout the literature: "a creative response is novel, good, and relevant." [56, xiii]. From a comfortable seat on board in 1966, Liam Hudson wrote:

'Creativity' [...] applies to all those qualities of which psychologists approve. And like so many other virtues [...] it is as difficult to disapprove of as to say what it means. As a topic for research, "creativity" is a bandwagon; one which all of us sufficiently hale and healthy have leapt athletically abroad. [57, p100-101].

But why? What are the reasons for the astonishing success of the Creativity bandwagon, which continues to gain speed, and has left in its wake a whole set of often quite unrelated "creative industries" (media, advertising, TV, film, design, games)? Even

[6] Vygotsky had a similar view: "just as electricity is equally present in a storm with deafening thunder and blinding lightning and in the operation of a pocket flashlight, in the same way, creativity is present, in actuality, not only when great historical works are born but also whenever a person imagines, combines, alters, and creates something new, no matter how small a drop in the bucket this new thing appears compared to the works of geniuses."

banking is given the epithet "creative" without a trace of irony, as well as the great entrepreneurs, led by Richard Branson. Here are just a few of the possible reasons for this remarkable juggernaut.

A. It is held together by the scientific armour of Factor Analysis, a way of constructing smooth curves from the uncertain data of questionnaires.
B. Protected by this show of rigour, it was able to break away from the aridities of Behaviourism, which had given Psychology its needed scientific respectability but had bored students for years.
C. The giants of Humanistic Psychology got on board, each with a mouth-watering trade mark to draw students to Creativity 101: Carl Rogers' self-actualization in 1954 [58], Csikszentmihalyi's flow in 1975 [59], and Maslow's peak experiences in 1968 [60]. Charles Tart was there with altered states of consciousness in 1969 [61], and Frank Barron, veteran of LSD experiments in 1963 . And even Buddhism, offering an endless stream of books with titles beginning "Zen and Art of . . . ." to say nothing of Kabat-Zinn's introduction mindfulness as an essential component of creativity in 1990 [63]. It all added much needed glamour to Psychology.
D. Artificial Intelligence hitched a lift. As early as 1958 Newell et al [64], had raised the problem of creativity for computers and described a programme on ILLIAC that composed music. Computational creativity has progressed independently (there are remarkably few cross references between the two disciplines) but in parallel with Psychology's version, and has probably added a further bit of hard-nosed scientific respectability to the whole endeavour.
E. Last but not least, there has been massive funding from military and industry. As Guilford wrote in 1959, soon after the launch of Sputnik by the USSR *"The preservation of our way of life and our future security depend upon our most important national resources: our intellectual abilities and, more particularly, our creative abilities. It is time, then, that we learn all we can about those resources"* [66, p469]. The economy and safety of the West is thought to depend on the practical benefits of making things that work, from nuclear weapons to the stylish artefacts of Steve Jobs, and the secret is creativity.

3. Creative Feedback

But in the midst of all this razzmatazz, there was a quiet Deweyan revolution. Some of it took place on the bandwagon itself, where there are researchers who stress that Creativity is an everyday matter, and that we all possess it in our capacity for flow and mindfulness. More recently there are those who have turned away from creativity with a capital C, and looked at how a more modest Deweyan creative intelligence can be encouraged throughout education [67,68,69]. Dewey believed that creative intelligence is necessary for democracy to prosper, and it is fostered by what we call creative feedback.

This is the goal of MusicCircle Software project at Goldsmiths; to design an online environment to support communities of creative feedback for learning to play music [71]. It includes the ability to upload performances, share them with others, and then seek and provide creative feedback. It is developed through a process of participatory design, working with students and other users to ensure we build what people want.

Through systems such as ours perhaps we can begin to reconcile the conflicting demands of Michael Gove and Ken Robinson through evidencing clearly how learning takes place through creative feedback.

In order to understand how to design learning environments, we now set out to characterise creative feedback in more detail. We do so by describing its qualities along a number of dimensions drawing both upon our historical analysis and our combined backgrounds: teaching, programme development and management in higher education; performance and composition in music; design and implementation in software; and mindfulness and psychotherapy in practice. These qualities of creative feedback are offered in hope of receiving creative feedback to inspire the next steps.

1. CF is social. It comes from one social agent who has perceived the feedback object in some way (whether that is an output or a process of an individual) to another (the originator of the feedback object). Note this definition does not preclude students giving creative feedback to their own work.
2. CF is mindful. This incorporates at least two aspects: a) that the person giving the CF is aware of the cultural and individual context of the receiver (such as an understanding of the individual's artistic or scientific goals/methods/audiences etc.); and b) that individuals are aware of any personal judgments that are being made and can articulate these if required.
3. CF contains a degree of community awareness that: a) CF embodies an awareness of what creative feedback has occurred previously but also that it features as part of a complex and developing system; and b) giving and receiving CF should be embraced equally for the community to sustain itself. It would be difficult for communities to thrive if everyone wanted to give more CF than they wanted to receive of course. CF creates a self-sustaining self-organising system where flexibility and robustness need to be balanced. Whilst each learner may have more or less knowledge about what is required to maintain such a system it is clear that it can only exist if individuals in the learning environment actively encourages engagement in CF.
4. CF is clear, the language used being unambiguous and terms used mutually understood.
5. CF is democratic. Being a tutor or student bestows no special right to giving or receiving CF (though of course one might hope that tutors have more experience and skills in giving it).
6. CF is challenging. Underpinning any creative partnership is the notion of the challenge that the each brings to the other. CF that provides the right level of challenge is arguably the most sought after feedback. To do so involves "skill in means", a Buddhist concept meaning that feedback is geared to the level and character of the student, and is always open to the student's needs.
7. CF incorporates generosity of spirit and compassion. It is an act of giving and enabling, itself an essential aspect of skill in means.
8. CF is always open to discussion and further explanation.
9. CF is comparative rather than absolute. No absolute judgment about a feedback object can be made. Comparisons (explicit or implicit) of the feedback object to other existing objects is a mindful tactic in many cases and involves skill in means. (For example, CF to a jazz piano student from a tutor could simply say how close

the student's playing is to another well-known jazz pianist and how they may want to take a listen.) We believe the key to successful education is about providing the right kinds of environments where skills in creative feedback can develop. The role of technology is both to build new kinds of learning environments but critically to start to evidence how the creative feedback ability is correlated with learning and artistic development more generally. This may have ramifications for the way in which we think about structuring learning in schools, universities and any other kind of learning community.

4. Concluding thoughts

We are designing a new technology at Goldsmiths called Music Circle as part of a European Project (Practice and Performance Analysis Inspiring Social Education) through the Technology-Enhanced Learning Programme. [71] It is designed to allow students to upload and share performances and compositions within learning communities and then by inviting feedback from others. In order to identify the kind of feedback we wish to encourage in our system (which currently operates in a blended learning context at Goldsmiths) we have identified the term "creative feedback" which embodies a range of characteristics including clarity, mindfulness, generosity, challenge and democracy.

At the heart of the motivation for designing this system is the idea that students can learn a huge amount from the creative feedback given by others. Not only that, but that the students can develop their own abilities as musicians through the ability to give creative feedback to others. And there is little doubt that the ability to receive feedback well, to depersonalise it as much as possible and respond to it appropriately, will stand students in good stead for the world of professional musicianship. Moreover, outside the professional music world, employers will be seeking students who have the skills to work in communities that have skills in giving and receiving creative feedback. Indeed one can easily imagine a world where an employer is much more interested in the way in which a student has contributed to and benefitted from being in a community. So our manifesto and agenda for change may result in students leaving universities not with a transcript of module marks but with a detailed account of their sustained engagement with creative feedback in a community of learners.

As part of the design of the system, we are designing "creative feedback agents" that are software systems that can start to provide some aspects of creative feedback on uploaded performances and compositions. With the development of techniques from audio analysis, gesture analysis, and style analysis combined with building models of learners we are looking to build systems that can start to embody some of the CF characteristics we have identified in this paper. What is important to us is that the design of our software is underpinned by a strong educational philosophy that comes from an understanding of the historical precedents and discoveries of many before us. We want to move away from the idea that technologies are designed and built by technologists and we embrace a multi-disciplinary approach where learners, educators, designers, sociologists, philosophers, historians, psychologists and computer scientists come together to build systems but with a clear understanding of the work that has come before. Perhaps more than anything this paper is a call to arms to revive and embed a Deweyian educational philosophy that can now be both supported and evidenced through technology.

5. Acknowledgements

Our thanks to Goldsmiths, Harry Brenton, Roger Burrows, Rosie Shepperd, Matthew Yee-King, Francois Pachet, Jon McCormack, Andreu Grimalt-Reynes, Melly, Maisie and Maureen Still, Sarah Khan, Jonathan James, Chris Kiefer, Carles Sierra and Robert Zimmer. This research was supported by the FP7 Technology Enhanced Learning Program Project: Practice and Performance Analysis Inspiring Social Education (PRAISE) which includes Goldsmiths, Sony Computer Science Laboratories in Paris, the Institute of Artificial Intelligence in Barcelona and VUB, Brussels.

References

[1] d'Inverno, M. and M. Luck (2012). *Creativity through Autonomy and Interaction.* Cognitive Science, 4(3): 332-346.
[2] Sudnow, D. (1978). *Ways of the Hand.* London, Rooutledge & Kegan Paul.
[3] Mead, G. H. (1964). *Selected Writings.* Chicago, University of Chicago Press.
[4] Coffield, F. and B. Williamson (2011). *From Exam Factories to Communities of Discovery: The democratic route.* London, Institute of Education.
[5] Lewis, C. C. (1995). *Educating Hearts and Minds.* Cambridge, Cambridge University Press.
[6] Coffield, F. (2007). *Running ever faster down the wrong road: An alternative future for education and skills.* London, Institute of Education.
[7] Tomlinson, S. (1997). *Edward Lee Thorndike and John Dewey on the Science of Education.* Oxford Review of Education 23(3): 365-383.
[8] Thorndike, E. L. (1898). *Animal Intelligence: an experimental study of the associative processes in animals.* Psychological Review Monograph, No 8.
[9] Thorndike, E. L. (1911). *Animal Intelligence.* New York, Macmillan.
[10] Woodworth, R. S. (1918). *Dynamic Psychology.* New York, Columbia University Press.
[11] Bechtel, W. and A. Abrahamsen (1991). *Connectionism and the Mind: an introduction to parallel processing in networks.* Oxford, Blackwell.
[12] Thorndike, E. L. (1904). *An Introduction to the Theory of Mental and Social Measurements.* New York, The Science Press.
[13] Dewey, J. (1889). *Review of Natural Inheritance by Francis Galton.* Publications of the American Statistical Association 1(7): 331-334.
[14] Manicas, P.T. (2002). *John Dewey and American Psychology.* Journal for the Theory of Social Behaviour 33(2): 267-294.
[15] Dewey, J. (1938 (1963)). *Experience and Education.* New York, Collier Books.
[16] Lagemann, E. C. (1989). *The plural worlds of educational research* History of Education Quarterly 29(2): 185-214.
[17] Mchitarjan, I. (2000). *John Dewey and the development of education in Russia.* Studies in Philosophy and Education 19(1-2): 109-131.
[18] Vygotsky, L. (1978). *Mind in Society.* Cambridge, MS, Harvard University Press
[19] Tomlinson, S. (1997). *Edward Lee Thorndike and John Dewey on the Science of Education.* Oxford Review of Education 23(3): 365-383.
[20] Bruner, J. (1966). *Toward a Theory of Instruction.* Cambridge, MA, Harvard University Press.
[21] Wood, H. and D. Wood (1999). *Help seeking, learning and contingent tutoring.* Computers & Education 33: 153-169.
[22] Young, M. F. D. (2008). *Bringing Knowledge Back In: from social constructivism to social realism in the sociology of education.* Abingdon, Oxfordshire, Routledge.
[23] Wiener, N. (1948). *Cybernetics.* New York, Wiley.
[24] Miller, G. A., et al. (1960). *Plans and the Structure of Behavior.* New York, Holt, Rinehart and Winston.
[25] Chomsky, N. (1957). *Syntactic Structures.* The Hague, Mouton.
[26] Miller, G. A., et al. (1960). *Plans and the Structure of Behavior. p33, 43* New York, Holt, Rinehart and Winston.

[27] Mead, G. H., Ed. (1982). *The Individual and the Social Self.* Chicago, University of Chicago Press.
[28] Gibson, J. J. (1966). *The Senses considered as Perceptual Systems.* Boston, Houghton-Mifflin.
[29] Brooks, R. (1991). *Intelligence without representation.* Artificial Intelligence 47: 139-159.
[30] Lave, J. (1988). *Cognition in Practice.* Cambridge, Cambridge University Press.
[31] Robbins, P. and M. Aydede, Eds. (2009). *Situated Cognition.* Cambridge, Cambridge University Press.
[32] Thurstone, L. L. (1923). *The Stimulus-Response Fallacy in Psychology.* Psychological Review 30: 354-369
[33] Thurstone, L. L. (1927). *A law of comparative judgement.* Psychological Review 34(4): 278-286
[34] Cronbach, L. J. and P. E. Meehl (1955). *Construct validity in psychological tests.* Psychological Bulletin 52: 281-302.
[35] Cronbach, L. J. (1957). *The two disciplines of scientific psychology.* American Psychologist. 12(11): 671-684.
[36] Scriven, M. (1967). R. W. R. M. Tyler, R. M. Gagn and M. Scriven.*The methodology of evaluation. Perspectives of Curriculum Evaluation.* Chicago, Rand McNally.
[37] Shute, V. J. (2008). *Focus on Formative Feedback.* Review of Educational Research 78(1): 153-189.
[38] Arieti, S. (1976). *Creativity: The Magic Synthesis.* New York, Basic Books.
[39] Bront, E. (1995). *Wuthering Heights.* London, Penguin.
[40] Cardoso, A., et al. (2000). *An Architecture for hybrid creative reasoning. Soft Computing in Case Based Reasoning.* S. K. Pal, T. S. Dillon and D. S. Yeung, Springer: 147-178.
[41] Bergson, H. (1911). *Creative Evolution.* London, Macmillan.
[42] Dewey, J. et al (1917). *Creative intelligence. p23.* New York, Henry Holt.
[43] Dewey, J. (1934 (1980)). *Art as Experience.* New York, Perigree Books.
[44] Meyer, S. (2005). *Introduction: Whitehead Now.* Configurations 13(1): 1-33.
[45] Whitehead, A. N. (1929 (1978)). *Process and Reality. p208.* New York, The Free Press.
[46] Dewey, J. (1937). *Whitehead's Philosophy.* The Philosophical Review 46(2): 170-177.
[47] Mead, G. H. (1936). *Movements of Thought in the Nineteenth Century. p210.* Chicago, University of Chicago Press.
[48] Mead, G. H. (1934). *Mind, Self and Society.* Chicago, University of Chicago Press.
[49] Dewey, J. (1896). *The reflex arc concept in psychology.* Psychological Review 3: 357-370.
[50] Guilford, J. P. (1950). *Creativity.* American Psychologist. 5(9): 444-454.
[51] Spearman, C. (1904). *General Intelligence* Objectively Determined and Measured." The American Journal of Psychology 15(2): 201-292.
[52] Spearman, C. (1930). *The Creative Mind* London, Nisbet & Co.
[53] Torrance, E. P. (1962). *Guiding Creative Talent.* New York, Prentice-Hall.
[54] Duncker, K. (1945). *On problem solving.* Psychological Monographs 58(5, Whole No. 270).
[55] Mednick, M. T., et al. (1964). *Incubation of creative performance and specific associative priming.* Journal of Abnormal and Social Psychology 69: 84-88.
[56] Kaufman, J. C. and R. J. Sternberg, Eds. (2010). *The Cambridge Handbook of Creativity. xiii.* Cambridge, Cambridge University Press.
[57] Hudson, L. (1966). *Contrary Imaginations.* London, Methuen.
[58] Rogers, C. R. (1954). *Towards a theory of creativity.* ETC: A Review of General Semantics 11: 249-260.
[59] Csikszentmihalyi, M. (1975). *Beyond Boredom and Anxiety: Experiencing Flow in Work and Play.* San Francisco Jossey-Bass.
[60] Maslow, A. H. (1968). *Toward a Psychology of Being* New York, Wiley.
[61] Tart, C., Ed. (1969). *Altered States of Consciousness.* New York, Wiley.
[62] Barron, F. (1963). *Creativity and Psychological Health.* Oxford, Van Nostrand.
[63] Kabat-Zinn, J. (1990). *Full Catastrophe Living* New York, Delacorte.
[64] Newell, A., et al. (1958). *The processes of creative thinking.* Presented before a symposium at the University of Colorado, May 14, 1958.
[65] Vygotsky, L. (2004). *Imagination and creativity in childhood.* Journal of Russian and East European Psychology 42(1): 7-97.
[66] Guilford, J. P. (1959). *Three faces of intellect. p469.* American Psychologist. 14(8): 469-479.
[67] Claxton, G., et al. (2006). *Cultivating creative mentalities: a framework for education.* Thinking Skills and Creativity 1(2): 57-61.
[68] Gauntlett, D. (2011). *Making is Connecting: The Social Meaning of Creativity, from DIY and Knitting to YouTube and Web 2.0 London* Polity Press.

[69] Johnston, J. S. (2006). *Inquiry and Education: John Dewey and the quest for democracy.* Albany.
[70] Steels, L. (2015) The coming of (social) MOOCs. Chapter 1. In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.
[71] Yee-King,M. M. Krivenski, H. Benton, A. Grimalt-Reynes, and M. d'Inverno (2015) Music circle: an educational social machine. Chapter 7. In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

This page intentionally left blank

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-59

Chapter 4.
Social Flow in Social MOOCs

Luc Steels
ICREA, Institut de Biologia Evolutiva
Universitat Pompeu Fabra and CSIC

Abstract. Flow theory is a way to explain how humans can be self-motivated and reach a state of high focus and intense, very effective learning. Usually this theory is merely descriptive but recently it has also been operationalized and used as the basis for building autonomous agents. This paper examines how such an operationalization can be incorporated in computer-supported learning environments such as MOOCs. It also expands the notion of flow to take into account 'social flow' occurring in a group of learners, such as a sports team or a small Jazz ensemble. We discuss how this kind of social flow can be induced, what the benefits are, and how it is relevant for building learning communities through web and social media.

Keywords. Flow, Social Flow, Autotelic principle, MOOCs, online learning, distance-education, social MOOCs.

1. Introduction

What is the best way to teach and learn? Whole libraries have been filled with theoretical treatises, many educational experiments have been performed (e.g. by Freynet, Malaguzzi, Steiner, Froebel, etc.), and a wide variety of national policies have been formulated and implemented. This rich landscape of educational theory and practice can be organised according to several dimensions. One important distinction is in terms of structured teaching versus open-ended learning environments, and a second one in terms of external versus internal motivation.

Structured vs. open-ended learning

The *structured teaching* approach assumes that teaching must rely on a carefully planned process that is then followed step-by-step by the learner. It is often practiced in state-wide education systems which are based on top-down fully planned national educational curricula where all learners of a certain age go through the same steps. The curricula define what has to be learned, which handbooks must be followed, how courses are to be prepared by the teachers, which exercises must be carried out. The outcome is carefully monitored through standardised, often nationally administered exams. In this scheme, teachers are often reduced to executioners of centrally designed plans and learners are viewed as uniform absorbers of knowledge. The teacher is also the source of authority, challenging the students and handing out reward and punishment. Although structured learning usually takes place in a classroom, the learner is individually evaluated and supposed to master the material by personal study.

Structured teaching has dominated the European educational landscape since the start of national curricula in the 19th century, and it is practiced world-wide. It is also widely used for continued education. Structured teaching works well for certain domains of knowledge. It has important advantages in terms of quality control, uniform training of teachers, and job recruitment - because it is clear what a diploma stands for. On the other hand, it is not very well adapted to the current generation of restless students, which are not only confronted with a massive amount of information to learn from, but also with a bombardment of multi-media materials, delivered through social media such as Facebook or Twitter, that encourage very short and rapidly shifting attention spans. It is also not very well adapted to types of expertise that require a great dosis of creativity, such as Jazz improvisation or advanced programming. Nevertheless, today's MOOCs mostly follow such a highly structured pedagogy, with a fixed lesson-plan that presents small chunks of knowledge or methods to students, and provides regular tests to check whether the student has mastered the material.

An *open-ended learning* approach is regularly proposed as an alternative that would lead to greater motivation and direct participation of students, and hence to a more enjoyable learning experience and continued curiosity. In this approach, the students are presented with a rich learning environment and a library of challenges they can tackle. They must seek out themselves the knowledge and resources needed to cope with a challenge and they often need to collaborate with others to solve tasks.

An example of this approach has been practiced the past decades in the Reggio Emilia schoolsystem. Young children work in small teams on projects and have a wide range of tools and materials at their disposal. The projects require learning many skills (drawing, writing, calculation) but these skills are acquired within the context of the project, and teachers act as organizers and tutors [21]. The LOGO environment designed by Seymour Papert is another example implementing the same philosophy, later used as the foundation for the Lego Mindstorms electronics kit. LOGO provides programming primitives and an intuitive model of computation, straightforward enough to make sense to children [10]. Although LOGO can be used in a traditional classroom setting with a rigid lesson plan, Papert's original goal was to create a learning environment in which children could autonomously seek out challenges and gradually discover solutions and build up skill.

Hackatons work along the same principles. There is a target domain to be learned (for example, programming apps for a smart phone) and students work together in groups with a more knowledgable guide on hand but without a strict lesson plan. Within the domain of music, a small Jazz band is another good example of an open-ended learning environment [7], [9]. Band members learn from each other and by playing with more experienced players. They are challenged by performance before an audience.

A final example is a well-functioning team of ph.D students working together under the guidance of a professor. They tackle together challenging problems, exchange knowledge, acquire new skills together or seek outside knowledge.

Open-ended learning environments are 'learner-centered'. Each student carves out his or her own path and cooperation with others plays a very important role. There is usually a tutor but that is not even necessary. The advantage is that students are more motivated because they learn by doing and understand the relevance of the study material. It works best with good students that can move at their own pace, and have already mastered basic study material. For example, you cannot participate in a Jazz ensemble without prior music theory and competence on the instrument. On the other hand, external quality control becomes more difficult, the role of the teacher has to be much more proactive, and students may show big knowledge gaps because they may not have encountered a basic method or technique yet - although they have acquired the skills to fill those gaps autonomously.

Several researchers are currently exploring in how far open-ended learner-centered environments can be created using the web and social media, characteristic of MOOCs [30], [19]. They thus try to build *social MOOCs* [27], going back to the original idea of a MOOC, namely a platform to foster communities of learners [23].

These social MOOCs do not follow a strict lesson plan but provide a framework with the following basic components:

1. Materials and challenges to stimulate students. For example, a library of Jazz standards against which students can practice their improvisation skills [17].
2. Ways to share engagement with these learning materials, for example, by uploading improvisations performed by the learner. [19].
3. Facilities for giving peer commentary. For example, by attaching praise or criticism to a short stretch of uploaded music. [30]

On top of these facilities various elements are added to enhance the learning experience, for example:

1. Apps that provide automatic feedback to the individual, e.g. analyzing whether classical compositional rules have been violated [2], or whether the improviser is playing 'on the beat' [19].
2. Mechanisms to track opinion dynamics, e.g. using techniques of sentiment analysis, or network structure (who is interacting with whom) [13].

External vs. Self-Motivation

Another related dimension for categorizing the landscape of learning and education is in terms of the mechanisms that are available for fostering motivation. The high drop-out rate of current MOOCs and the large student failure in traditional instructional teaching shows that motivation is a critical factor that must be addressed. Here an important distiction can be made between external versus internal or self-motivation.

External motivation is based on reward and punishment. The reward can take the form of praise, higher marks, recognition by teachers and peers, fame, promotion, monetary benefit. Punishment can be in the form of low grades, public shaming, lack of recognition, demotion, rejection by peers. Instructional teaching, particularly in its extreme forms proposed by Skinner, is entirely based on the assumption that reward and punishment has to lie at the heart of education, not only to shape the knowledge students have to learn - because it provides positive and negative examples - but also to motivate them.

Internal motivation is entirely intrinsic to the learner. Several possible motors for internal motivation have been proposed (see the review in: [1]). One is based on the notion of curiosity [16], whereby the learner is driven by the desire to increase his or her success in predicting how the world behaves. Another possible motor for internal motivation is based on the notion of flow [20]. Flow is a state where humans obtain a strong focus on a particular task, losing awareness of all aspects of the environment which are not relevant for the task, feel great enjoyment, and achieve a high sense of creativity [5]. Concentration can be maintained for a very long time and no physiological effort is felt (although heavy fatigue may set in later). One of the most important characteristics of the flow state, is that the individual becomes *autotelic*, where 'auto-' means "self" and 'telos' means "goal". There is no need for external rewards.

As pointed out first by Cskiszentmihalyi [4], flow is observed in children at play, in athletes, artisans, artists, scientists but just as well people carrying out mundane looking jobs. They all go to great length to excel and push the boundaries of their knowledge and skills. Practitioners are said to reach a state of flow or optimal experience which not only pushes forward their competence but also provides tremendous enjoyment so that they seek the same experience again in the future. Most importantly, they see their activity as meaningful and the hard effort required for learning hence becomes meaningful as well. There is no doubt that this state of flow is reached by many musicians, including amateur musicians, and that this is the main reason why they go to great lengths to master their instrument or engage in music theory, a rather dry subject in itself.

Most humans operate both with internal and external motivation, and some people are more apt to experience and hence seek flow than others. But there is a consensus that the high emphasis on external reward in the structured learning environments dominating current educational systems is leading to a large number of negative side effects, among them: shallow knowledge acquisition, just enough to pass grades, and not in the least, a lack of long-term motivation [12]. Many people report that their desire to learn an instrument was stifled in an early stage by teachers who applied a reward and punishment methodology - occasionally associated with physical violence. The experience of flow is encouraged in open-ended learning environments, indeed it is a crucial ingredient of such learning environments because it is the major organizing principle, determining when and why learners are autonomously seeking what knowledge or skill to acquire, and the major motivational principle, because the teacher (or an assessment board) is no longer the source of authority handling out punishment.

The present paper pursues the question how the concept of flow can help to build more exciting and effective MOOCs, tackling in particular the problem of student motivation and scaffolding of complexity in tackling knowledge. First we discuss models of flow in individuals so that we can identify the requirements of a learning environment inducing flow. Next we turn to flow in groups. I characterize the notion of 'social flow' and discuss its relevance for the creation of social MOOCs.

2. Models of flow

Most of the work on flow and self-motivation in psychology is of a descriptive nature, taking an observer's point of view and using questionnaires and experience-sampling methods [14]. The main application area is counseling towards achieving well-balanced personalities or more productive work places [22]. However there has also been some work on building artificial 'autotelic agents' which incorporate the kind of mechanisms assumed to be necessary for flow as computational components in (robotic) agents that learn autonomously by interaction with their environment and others [25] [26].

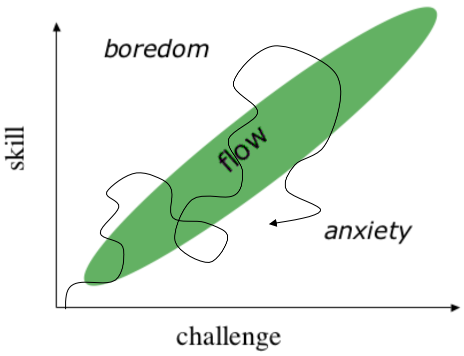

Figure 1. Flow diagram displaying development of a single individual. The x-axis shows the degree of challenge of tasks being considered and the y-axis the skill level of the individual. The flow state occurs when the two are balanced. When the challenge is too high for the skill, anxiety sets in. When the challenge is too low, the subject experiences boredom. During development a person navigates around the flow regime, regulating the challenge level of tasks he or she takes on while skills build up due to learning or decrease due to lack of practice.

2.1. *Descriptive models*

The key idea to come out of empirical research on optimal experience, already stated by Cskiszentmihalyi [4], is that there are a number of critical characteristics of flow. The first set of characteristics are preconditions for reaching flow:

1. There must be a clarity of goals. The individual has to be entirely focused on trying to achieve a particular challenge and the resources are available to carry out the necessary subtasks.
2. The goal is set by the individual, so that there is a sense of control. When the goals are beyond reach, the individual must be allowed to change goals, decreasing the challenge.
3. There must be clear feedback on how well the individual is doing, i.e. whether actions make significant steps towards the goal.
4. There must be a balance of challenge and skill, as captured in the famous flow diagram (see Figure 1). The task should fit well with available skills to avoid anxiety but also challenging enough to avoid boredom.

The second set of characteristics relate to the subjective experience while being in the flow state:

1. Awareness and action are merging, so that time seems to disappear.
2. There is a high concentration.
3. Solutions to problems come spontaneously and rapidly.
4. There is no concern or worry about anything else.
5. The activity is felt as intrinsically rewarding and therefore repeated.

More recent literature sees flow more as a process, in two ways. First, the flow state itself is dynamic: there are various preconditions, it takes a certain amount of time to get into it, and there is a phase of winding down with consequences for future behavior [11]. Second, the challenge/skill landscape is constantly shifting [26]. Skill builds up during the exercising of a task, so that a task which was a big challenge at some point may no longer be so and therefore may become boring. On the other hand, skill may decrease when not exercised, forcing the learner back to earlier challenge levels. This is a common experience of musicians, who have mastered a particular piece after a lot of hard work but then find that they have to go through a tedious (and less motivating) learning process again when they have not practised this piece for a while.

2.2. Autotelic Agents

Complementary to descriptive models, there have been efforts the past decade to develop agent-based models of flow [24]. 'Agent-based' means, that we think about the cognitive mechanisms, context, and social interactions that are required for the flow state to occur, and proposals are validated using computer simulations or experiments with robotic agents [25]. This objective requires a significant change in perspective because both challenge and skill are not parameters that are under direct control. It is often difficult for a learner to know in advance how challenging a particular task will be and skill cannot simply be increased at will, it requires practice, possibly with tasks that are less challenging.

Modeling is further complicated because a particular task usually requires a wide variety of components. For example, playing a particular piano piece fluently may be hampered by lack of skills in sight reading, unfamiliarity with the tonality of the piece, unusual chords or arpeggios, or simply lack of practice to move fingers fast enough. Each of these skills has its own requirements and a learner may have reached uneven levels of skills.

The Steels flow model [24] has been applied both to robot behavior acquisition [25] and language learning [26] and has recently been reimplemented by Cornudella et al. [3]. It assumes that an autotelic agent has a set of self-developing components, all necessary to solve a complex task within a particular situational context. The components have input-output relations with each other and may be organized in a hierarchy.

Each self-developing component has the following elements (see Figure 2):

- *Goal*: The component performs some mapping from inputs to outputs.
- *Knowledge*: The mapping is established through an algorithm (which may have critical parameters), a neural network, a set of production rules, or any other kind of computational device.
- *Feedback*: The component receives a feedback signal how well the mapping was established, coming from other components that make use of the result, or from an external source.
- *Learning*: Each component has a learning mechanism which is responsible for learning the knowledge needed for establishing the mapping. Any kind of learning mechanism can be employed.
- *Performance:* The component is able to track how well it has established a mapping.
- *Challenge level* The component is able to determine the level of challenge of the input and has at any point in time a particular challenge level as target. The challenge consists of a number of aspects and is represented as a feature vector with values between 0.0 and 1.0. Inputs whose challenge level surpass the challenge level set by the component are ignored, leading to overall failure in the task be-

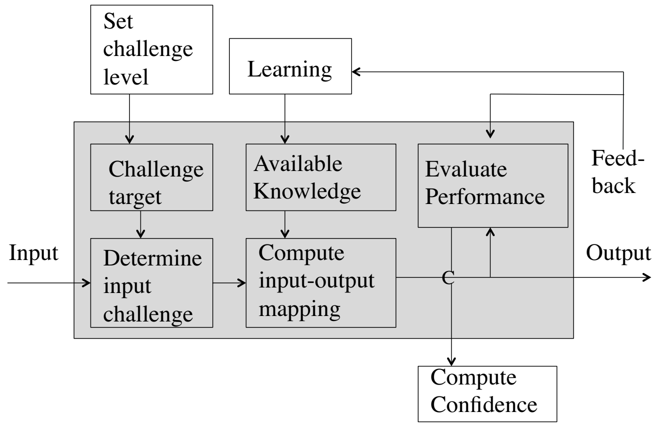

Figure 2. A self-developing component is able to compute a mapping between input and output, but it has additional machinery to determine whether the input is within the given targeted challenge, to learn the mapping and evaluate performance based on feedback, and compute confidence over time.

cause other components relying on it will receive no input. In principle a self-developing component should include an additional mechanism to learn which aspects of the input are important and how they can be calculated but this point has not been worked out further in the models developed so far.

- *Confidence*: The component has a way to track in how far it masters its goal. It is based on average performance over a window of time for a given challenge level.

The objective of a component is to reach steady performance for the targeted challenge level. Because usually some form of learning is required, it may take a while before this point is reached. When there is steady performance and hence high confidence, the challenge level can be increased. When feedback persistently generates a lot of failures, the challenge level is decreased to build first enough expertise to attempt a higher challenge level later.

Consider by way of example the task of sight reading and performing a piece of music. Human learners typically spend years before they fully master this task, gradually scaling up the difficulty with many hours of practice required. Suppose we want to build an autotelic agent that goes through the same developmental process. For example, imagine a robot that tries to ‘read’ a score using a camera and plays the notes on a saxophone. Two components are needed: the ‘sight reading component’ decides which note needs to be played and the ‘note production component’ produces the needed physical movements to play the note.

Let us just focus on the first component: deciding which note to be played. It needs the following elements:

- *Goal*: The input is a score and the output is a stream of instructions to the motor component that controls manipulation of the instrument.
- *Knowledge*: The component needs to visually recognize the pitch of each note on the staff, taking into account sharps or flats, the key, and the tuning of the instrument.

- *Feedback*: Given current technology, feedback can easily be provided by an automatic system that compares the produced note with the desired note, for example based on a rendering of the same note on a synthetizer.
- *Learning*: Associations need to be acquired between the visually perceived position of a note on a staff and the tone to be produced. Practise leads to more accurate and faster retrieval of this relation.
- *Performance*: Performance is determined by the speed with which the note is identified and whether it is the correct note when played.
- *Challenge level*: Challenge decomposes into a number of parameters, such as the distances between the notes (shorter distances are usually easier), how common the interval is (e.g. a third is usually easier than a sixth), the key (C-major with no sharps or flats is easier than C#-major which has seven sharps), the speed ('largo' is easier than 'presto').
- *Confidence*: This measure computes in how far the component masters its goal. It is based on average performance over a window of time for a given challenge level.

An autotelic agent either receives inputs from the environment without being able to control their challenge level, in which case he just ignores cases that cannot be handled or handles a situation only partly (e.g. an instrumentalist may skip ornamentations that require very fast finger movements). Alternatively, when an agent can control which situations he will handle (e.g. the instrumentalist can choose which scores to try) he can use a prior evaluation of the challenge level of the situation and choose those that fit with the targeted challenge level.

When there is a set of such self-developing components that are interrelated in the sense that one provides input to the next one, the agent is confronted with a multi-dimensional control problem through which he has to navigate, increasing or decreasing the challenge levels until stable performance is reached, and then climbing up by selected increases of the challenge levels of individual components. For example, the production of tone and the fingering of the instrument needs to be mastered at the same time as reading scores. Even if sight reading is fast enough, the instrumentalist might still not be able to fluently play the notes within the required tempo.

Occasionally the agent might get stuck in a loop that traps further growth: when challenge is decreased it leads to increased performance, but then increasing the challenge again leads (even after a period of learning) back to decreased performance. This signals that an additional or different component must be recruited for the task or that some other learning strategy is needed by one of the existing components.

An autotelic agent is striving towards two states: (i) to stay within the 'flow' regime, in the sense of having a consist performance which means that there is a balance between challenge levels and skill, and (ii) to achieve, in the longer term, a steady increase in challenge levels, while maintaining adequate performance. Note that this is clearly an intrinsic motivational system. There is no external agency that supplies the reward or scaffolds the external environmental conditions.

Optimizing a multi-dimensional system is known to be a notoriously difficult problem. It is faced, for example, by roboticists that have to control a humanoid robot with a large number of degrees of freedom. Many techniques exist for dealing with it. In the original Steels model [24], this issue is approached by introducing two different phases: a learning phase and a shake-up phase. During the learning phase, all challenge parameters are kept fixed and the agent is exercising knowledge or learning new knowledge. This phase lasts until there is a high confidence level. In the shake-up phase, the challenge parameters are changed. There are two possibilities:

- Performance is consistently low, leading to low confidence in task achievement (anxiety). In this case one or more challenge parameters need to be lowered, typically the ones that were increased most recently.
- Performance is consistently high, leading to high confidence in task achievement (boredom). In this case, some of the challenge parameters need to be increased.

It is quite often the case that performance is decreasing rapidly after the increase of challenge parameters, which requires the agent to take action sooner. The choice which challenge parameter in which component is to be changed is difficult and heuristics need to be employed to avoid combinatorial explosions. For example if component X depends on input from component Y, whereby Y has a high performance but X consistently fails, then one of the challenge parameters of Y can be decreased in order to allow X to catch up. Concretely, if sight reading itself is already smoothly working for complex scores but the instrument is not yet sufficiently mastered, then the targeted complexity of the scores should be of a lower challenge level, until the note production component catches up. There is no doubt that a good learner employs powerful heuristics like this, helping to choose when and how to simplify the problems he or she tackles, and when and for which aspect of the task the challenge level should be increased.

3. Social Flow

Most of the literature on flow focuses on how a single individual can reach a flow state and what beneficial effects this can have. But music and many other human activities often take place in a group and there is a widely shared belief that a group is more than the sum of its individual members. It has its own 'flow dynamics' influencing the selection of tasks and the availability of knowledge to share. A special case of social flow occurs when the 'group' is very small, consisting for example only of a learner and a tutor, or a child and her two parents.

There are two ways in which inter-individual dynamics enhances learning.

1. Some members in the group may have a higher skill level than others and are therefore able to solve a more challenging task, possibly explaining on the way how a solution can be reached. This can help to pull less knowledgable members towards a higher level of expertise and acts as a role model of what can be achieved by others. A tutor is a special case of this, but this situation can also occur in a group of peers.

Vygotsky's notion of the *zone of proximal development* conceptualizes this type of learning situation (Figure 3) [29]. He groups task challenges into three zones. The inner region are problems that the learner can solve comfortably, the middle zone are problems that the learner can solve when aided by a more knowledgable tutor or peer who can scaffold the problem by solving subproblems which the learner cannot solve yet, by drawing attention to aspects of the problem to which the learner is not yet sensitive, or by supplying information needed. The third zone is entirely out of reach.

Vygotsky's conceptualization is often brought in relation to the flow model (Figure 1) by equating the zone of proximal development with the flow region in which challenge

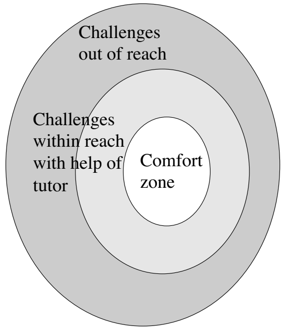

Figure 3. Vygotsky's conceptualization of learning guided by a tutor involves a zone of challenges between comfort (possibly leading to boredom), a zone with problems out of reach (possibly leading to anxiety), and a region where a tutor can scaffold the situation to allow the learner to discover solutions.

and skill is balanced, although this seems to be too much of a simplification. First of all, the flow model does not address tutoring situations but the optimal experience of an individual, and second, the zone of proximal development is not the region of comfort where there is a strict balance between challenge and skill, but rather the region where skill is insufficient, but thanks to scaffolding the learner moves towards the flow region.

2. The members of a group can often solve problems together which no single individual is able to solve. A group can then be viewed as a single autotelic agent with multiple components which solve specific subtasks. Because the members of the group have to cooperate, their challenge levels must be compatible. For example, a jazz ensemble has different players each contributing with their own instrument and having different functions (rhythm section, harmonic background, melodic lines). The problem is the same as between the components of a single agent, namely how to ensure that the challenge levels of the different components are compatible with each other. The group can only thrive and move up in challenge when the different players have roughly equal levels of competence. Humans spontaneously will lower complexity of their behavior for others, e.g. a mother uses 'motherese' to scale down the complexity of her language in order to create input that the child can master, a pianist may simplify the chord structure so that the beginning improviser hears more clearly the harmonic structure of the piece.

The autotelic model has been operationalized and used in various experiments in which robots acquire more complex behavior (which is the case of one autotelic agent against the environment that he needs to scaffold himself) or more complex language (in which you have multiple agents that have to coordinate their challenge levels to allow everybody to catch up).

An example of the kind of results obtained is shown in figure 4 from [26], which contains more details. The experiment concerns a population of 10 agents trying to create a common language. Challenge levels relate to the complexity of the meanings that agents try to express, moving from single predicates to relations. Initially they use a lexical strategy, that allows them to invent and learn words from others, but they are able to recruit a grammatical strategy if needed. Task success is here equal to communicative success in the language game. The size of the lexicon first overshoots and then agents align. Communicative success gets higher and higher and confidence builds.

Next challenge levels are being increased by agents (around interaction 4000) and there is an immediate significant drop in performance. Agents try at first to invent more

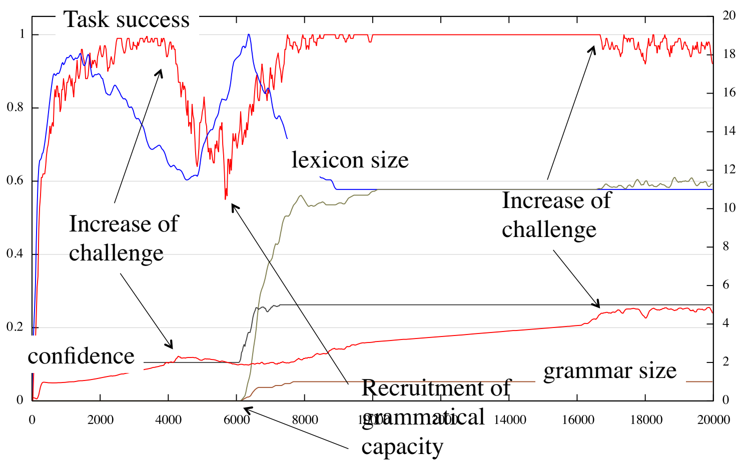

Figure 4. Experiment with social flow in a population of agents self-organizing a language. The x-axis shows the number of interactions between agents and the y-axis various observed measures: Task success, which is the running average of succeeded communications, lexicon size which gives the average size of the lexicon of all agents in the population, grammar size which gives the average number of constructions, average confidence of the agents, and some other measures not relevant for the present discussion.

words, but relief only comes when recruiting a grammatical strategy. As agents invent and share grammatical constructions, success goes back up, until a new cycle of challenge increase becomes feasible (around interaction 16000).

There has been a lot of work on operationalizing motivational theories based on reward and punishment, but these experiments show that it is also possible to operationalize autotelic principles and incorporate them in artificial agents. This field is in its infancy and many further experiments have to be done, for example to discover heuristics for navigating in a multi-component autotelic system. These operationalization will help to better understand the flow phenomenon and base learning environments on this theory of motivation.

4. Implications for MOOCs

The first experiments exploring flow for creating more exciting learning environments were conducted by François Pachet [18]. Pachet proposed to create ‘mirrors’ for a learner that reflect back his or her knowledge or skill, for example, playing back the chord sequences that the learner already knows, as a stimulus to then entrench existing knowledge or be a basis for building further on it. Pachet’s ‘Continuator’ is a music system that acquires the statistical properties of musical input introduced through a keyboard and then mirrors this with its own output which consists of variations on the learned patterns. This activity generates a lot of excitement, even in young children. The system learns continuously and therefore players try to elicit more complex behavior by becoming more sophisticated themselves. Pachet’s ‘flow machine’ (in the sense of machines that help to generate flow) tell us something about why humans become excited and what features a system needs to have in order to elicit excitement.

Another approach, illustrated by the (artificial) counterpoint tutor described later in this book [2], assumes that the learner is an autotelic agent and that the tutor must help to scaffold tasks and inputs and steer the learner through the search space of challenge levels. The (artificial) tutor achieves this by modeling the student at a very detailed level, and this allows him to come up with exercises that are within his or her Zone of Proximal Development.

More general, embedding autotelic principles in a (social) MOOC learning environment requires the following:

1. The system should generate clear goals in the form of tangible problems to be solved. In the case of music, these goals could be: generating an improvisational line on top of an existing accompaniment for a Jazz standard, writing a 4 part choral piece given a melody, interpret a Chopin prelude on the piano.
2. The learner must be able to select a goal to pursue, among a set of possible goals, and also the specific situation that is to be tackled (e.g. which melody will be harmonized). This goal should be compatible with his or her skill level. This implies that learners must either be given an indication of the challenge level required for a particular task (e.g. how hard is it to play along or improvise for a particular Jazz standard in the database) or they must be encouraged to build up skill for determining the challenge level themselves (e.g. by inspecting the score).
3. Learners must have a way to gauge how well they are doing, ideally in ways that do not kill the enjoyment of engaging in the task, in other words, not by separate tests, but by tracking performance during the execution of the task, if possible automatically. So feedback is of crucial importance and many of the mechanisms discussed in this book (e.g. testing the quality of a Jazz improvisation based on tracking how many notes are within the scale, what shifts occur with respect to the beats, whether melodic continuity is preserved, etc. [19]) are entirely relevant. In the case of a social MOOC, learners must be able to get an idea of the performance level of others so that they can choose peers with which to jointly solve a task (e.g. jointly engage in an improvisation).
4. The learning environment must not only track performance to regulate the challenge level of proposed tasks, but also make a model of the learner in order to present appropriate possible goals, related to the skill level of the learner.

Of course many school curricula and MOOCs already have such elements, such as continuous testing, scaffolding of course material, etc., and teachers will naturally attempt to personalize student tasks - although testing is usually uniform. The key difference is that a flow-based learning environment gives more control to the learner who can thus carve out his or her own individual trajectory to balance challenge and skill.

5. Conclusions

Flow theory is a welcome complement to the reward-and-punishment theory that underlies a lot of learning theories and teaching systems. It leads to greater long-term motivation, curiosity about the subject, and a more balanced feeling of well-being. More work is needed to understand the cognitive mechanisms and features of the task context that help the induction of flow, and these can be of great value in creating computer-based learning environments, including MOOCs that are more personalized to the individual student and present material and exercises adapted to move from the student's comfort zone to the Zone of Proximal Development. Particularly in the case of social MOOCs, flow theory points to certain characteristics that could in many cases be easily added to currently available systems, to make it easier for learners to have more enjoyable and effective learning experiences.

References

[1] Baldassarre, G. and Mirolli, M. (2013). Intrinsically motivated learning in natural and artificial systems. Springer Verlag, Berlin.

[2] Beuls, K. and J. Loeckx (2015) Steps towards Intelligent MOOCs. Chapter 9. In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

[3] Cornudella, M., P. Van Eecke, R. van Trijp (2015) How Intrinsic Motivation can Speed Up Language Emergence. Proceedings of the European Conference on Artificial Life 2015, pp. 571-578.

[4] Csikszentmihalyi, M. (1975) Beyond Boredom and Anxiety: Experiencing Flow in Work and Play. Cambridge University Press, Cambridge.

[5] Delle Fave, A. (2004). A feeling of wellbeing in learning and teaching. In M. Tokoro and L. Steels (Eds.), A learning zone of one's own (pp. 97-110). IOS Press, Amsterdam.

[6] Engeser, R. (2012) Advances in Flow Research. Springer Verlag, Berlin.

[7] d'Inverno, R. (2015) Teaching Jazz Improvisation. Chapter 5. In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

[8] d'Inverno, M. and A. Still (2015) Chapter 3. In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

[9] Jones, E. and H. Brenton (2015) Using social media to revive a lost apprenticeship model in Jazz education. Chapter 11. In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

[10] Kahn, K. (2015) A half century perspective on the role of computers in learning and teaching. Chapter 14. In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

[11] Kawabata, M., and C. Mallett (2011). Flow experience in physical activity: Examination of the internal structure of flow from a process-related perspective. Motivation and Emotion, 35,(4), 393-402.

[12] Kohn, A. (1999) Punished by Rewards. The Trouble with Gold Stars, Incentive Plans, A's, Praise, and Other Bribes. Houghton Mifflin, Boston.

[13] Loeckx, J. (2014) Mining for Evidence of Collaborative Learning in Question/Answering Systems. Workshop on Feedback from Multimodal Interactions in Learning Management Systems, 7th International Conference on Educational Data Mining (EDM 2014)

[14] Moneta, G. (2015) Flow in work as a function of trait intrinsic motivation, opportunity for creativity in the job, and work engagement. In S. Albrecht (Ed.), The Handbook of employee engagement: Perspectives, issues, research and practice (pp. 262-269). Edward-Elgar Publishing House.

[15] O'Neill, S. and G. McPherson (2012) Motivation. In: Parncutt, R. and G. McPherson (Eds) (2012) The Science and Psychology of Music Performance: Creative Strategies for Teaching and Learning Hardcover

[16] PY Oudeyer, F Kaplan, V. Hafner. (2007) Intrinsic motivation systems for autonomous mental development. IEEE Transactions on Evolutionary Computation. 11 (2)

[17] Pachet, F., Suzda, J. and Martn, D. (2013) A Comprehensive Online Database of Machine-Readable Lead Sheets for Jazz Standards. Proceedings of ISMIR 2013, pages 275-280, Curitiba (Brazil).

[18] Pachet, F. (2004) On the Design of a Musical Flow Machine. In: Tokoro, M. and L. Steels (eds.) A learning zone of one's own. IOS Press, Amsterdam.

[19] Ramona, M., F. Pachet and S. Gorlow (2015) Chapter 8. In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.

[20] Rheinberg, F. (2008). Intrinsic motivation and flow-experience. In H. Heckhausen and J. Heckhausen (Eds.), Motivation and action (pp. 323-348). Cambridge: Cambridge University Press.

[21] Rinaldi, C. (2003) The Joys of Pre-school learning. In: Tokoro, M. and L. Steels (eds.) The Future of Learning. Issues and Prospects. IOS Press, Amsterdam. p. 57-78

[22] Ryan, R. and E. Deci (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. American Psychologist, 55, 68-78.

[23] Siemens, G. (2005) Connectivism: A Learning Theory for the Digital Age. International Journal of Instructional Technology and Distance Learning, 2(1).

[24] Steels, L. (2004) Autotelic Agents: Robots that have fun. In: Hafner, V., F. Iida, R. Pfeifer and L.Steels (eds.) (2004) Embodied AI. Proceedings of the Dagstuhl Seminar. Lecture Notes in Computer Science. Springer Verlag, Berlin

[25] Steels, L. (2004) The architecture of flow. Chapter 7. In: Tokoro, M. and L. Steels (2004) A learning zone of one's own. IOS Press, Amsterdam.

[26] Steels, L. (2007) Scaffolding language emergence using the autotelic principle. IEEE Symposium on Artificial Life. ALIFE '07. pp. 325 - 332.

[27] Steels, L. (2015) The coming of (social) MOOCs. Chapter 1. In: Steels, L. (ed.) (2015) Learning about music in the age of (social) MOOCs. IOS Press, Amsterdam.

[28] Steels, L. and P. Wellens (2007) Scaffolding Language Emergence Using the Autotelic Principle. IEEE Alife Conference, Hawaii.

[29] Vygotsky, L. (1978) Mind in Society: Development of Higher Psychological Processes. Harvard University Press.

[30] Yee-King,M. M. Krivenski, H. Benton, A. Grimalt-Reynes, and M. d'Inverno (2015) Music circle: an educational social machine. Chapter 7. In: Steels, L. (ed.) (2015) Learning about music in the age of (social) MOOCs. IOS Press, Amsterdam.

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-73

Chapter 5. Teaching Jazz Improvisation: A Personal Experience

Ray D'INVERNO
University of Southampton

Abstract.
This short note gives a personal account of how I became an establish Jazz pianist. It emphasizes the importance of personal study and playing in a small group from the very beginning. I also describe various routes to improvisation, which I used as the basis for a MOOC on piano Jazz improvisation, currently in an experimental stage of development.

Keywords. Jazz improvisation, music learning, MOOCs for music.

I was classically trained as a pianist from an early age. When I was four years old, my mother died and, since she was a pianist, my father sent my brother, who was six years older than me, for piano lessons in her memory. My brother started off with a piano teacher but she eventually asked my father to withdraw him because he was not practising between lessons and therefore not making any progress. Later on, I went to the same piano teacher and, though initially I was making good progress, I eventually lost interest as well. The teacher again suggested to my father that he should withdraw me but, to his eternal credit, he refused. My lack of interest continued until one day I attended a piano concert in which I heard Christopher Sinding's "Rustle of Spring" played and it had a great impact on me. Even though my piano teacher considered the piece technically too difficult, I was resolved to learn how to play it and in the process I regained interest in the piano lessons.

I heard jazz for the first time when I was fifteen. I was listening to the radio one day when I heard the Hampton Hawes Trio playing "Gypsy in my Soul". I realised immediately that this was my music. From that moment on, I focused on trying to teach myself jazz as best I could bit in an era when jazz study materials were virtually non-existent. I did eventually find someone at my youth club who lent me a Dave Brubeck Quartet album. It was an EP - Extended Play - which, at that time, I had no idea of what it meant. I took it home, wound up my old 78 rpm gramophone, put the EP on it . . . and the needle scratched right across the record, destroying one side of it. So I then saved up for a Dansette record player that could play EPs and LPs and, in particular, the flip side of the Brubeck EP. The breakthrough for me came when I was lent an LP (Long Play) record of Eroll Garner. His style consisted of playing chords in the left hand, typically four-to-the-bar, and adopting that approach I had a rhythm section. I could then play the tune against this backing and improvise over the top of it. But what I really wanted to do was form a piano trio like the one I had first heard on the radio. In the end, I purchased a double-bass and a drum kit and first taught myself to play bass and drums. I then taught my best friend to play bass and another friend (whose name was, ironically, Dave Drum) to play the drums. I now had a trio. We were still only schoolboys but took every opportunity we could to play. We were were playing.

After a year or so of playing then together with some friends from school, we opened a jazz club in Bushey, North London - the BMJC - Bushey Modern Jazz Club, with my trio as the resident trio. We also invited established jazz musicians from London to play with us even though we knew so little and were so inexperienced. Playing alongside such excellent musicians really speeded up our learning process. After a while, I went up to Oxford University to read for a degree in mathematics and so was forced to leave the club behind. I ended up having to close it down in the Christmas vacation because the friends who had been left behind to run it in my absence had drunken all the profits. I even had to pay one of the last groups which had been booked, the Tubby Hayes Quintet, out of my university grant.

At Oxford, I got involved in the cabaret scene and I became musical director for the "Oxford Etceteras", It was in the year that included Michael Palin and Terry Jones, among others, who went on to become famous as members of "Monty Python". I also auditioned for an American musician called Eric Gutt who had been awarded a scholarship to Juilliard, presented by Dizzy Gillespie, for his trumpet playing. Eric did not take up the scholarship but instead studied for a PhD in philosophy in London University. However he chose to live in Oxford and advertised for musicians to form a band. I went to the audition but he did not like my piano playing style which was vey influenced by Oscar Peterson at the time, whereas his taste was more for the early bop pianists. However, I had very recently bought a flute and so in the event he formed a trio with him on piano, myself doubling on flute and drums, and a bass player named Richard Johnston. Eric also encouraged me to compose for the trio. He was a real original and between us I think we invented "free jazz" long before it became an established art form. We would sometimes get into such a “wacky” mood that the three of us would end up dancing and singing on top of the grand piano we used for rehearsals. Eric then became so influenced by a rather obscure philosopher called Stubbes that he changed his surname by deedpoll. We then became the "Eric Stubbes Trio". I think that in the end he obtained a temporary teaching post in philosophy at Aston University, but I never heard from him again.

After a while, I bought a record of the jazz pianist Bill Evans, and it had such an impact on me that I stopped playing in public for around six years. The music of Bill Evans really affected me: the style and approach was simply perfection to my ears. I thought I would never be able to play music like he did. I then only listened to albums of Evans and after a while I even stopped doing that because it had become so painful to do so. But having listed only to Evans throughout all this period it was perhaps not surprising that when I played I sounded very much like him. I had simply absorbed his style by osmosis. One day a colleague from University managed to get me to play reluctantly in public again in the interval spot in a jazz club in Southampton. The performance was pretty awful, but the organisers ended up inviting me to go to the monthly jam sessions which were held in a suite under the Town Hall of Southampton. I asked the compere Tim Colwell if I could play some piano and he asked me in what style. When I said "a bit like Bill Evans" he laughed and put me on first with an unknown bass player and an unknown drummer. But when he heard me play his face dropped, he fired his regular pianist on the spot and drafted me into his group "Tim Colwell's Jazz Friends". From then on I became involved in the jazz world on a professional basis playing in clubs and festivals and giving broadcasts in this country and overseas. Although fully committed to my day job as a university mathematics lecturer I was very busy as a jazz musician playing as many as 250 gigs in one year.

I then became involved in jazz education running a weekly jazz workshop for the the Adult and Continuing Education Department in Southampton University for some fifteen years. I was also heavily involved in church music running a very successful choir and instrumentalists. At one period the pianist, a wonderful musician, was the head of the Music Department in the university. Moreover, the then Dean of the Mathematics Faculty also had interests outside of mathematics. Through them I was allowed to teach in the Music Department for three years introducing a second and third year undergraduate course in jazz studies. I was replaced by an American jazz musician and jazz studies continues to this day. Indeed, I believe it is now possible to undertake some jazz studies in all three years of an undergraduate music degree and take a pathway in jazz. Southampton was the first provisional university to introduce this development and I believe that others have followed it.

Returning to the jazz workshop, it was originally a liberal arts programme that people took to learn a new skill. The workshops were held for three hours, one evening a week for a whole academic year. The aim of my programme was to teach jazz to students with no previous knowledge. The workshop included piano students who usually had a background in playing the piano but had no or little knowledge of jazz piano. The workshop proved to be very successful and some of my students even went on to professional jazz careers. Success in the course was of course largely to do with the efforts, practise and dedication of the student. Towards the end of my involvement in the workshop, the course ceased to be a liberal arts course and became credit-bearing. So the credits accrued could be used for university entrance or even towards a degree. The course now had assessment tasks and involved a final performance in the university concert hall with a professional rhythm section. The assessments were validated by the Guildhall School of Music and Drama. Not surprisingly, this fundamentally changed the nature and dynamics of the whole course and this, in the end, lead me to discontinue my involvement and I passed it on to a colleague.

The workshop included students with a wide spectrum of age, ability and background. In contrast, the music department course had students of much the same age and usually a schooled classical music background, but for whom jazz was a new area. The music course had around 20 or so students and it had a very practical approach, with the first part of the class devoted to theory and the rest for practise and performance. Not only did the students clearly enjoy their involvement but they gave the course in a final questionnaire the highest rating I have ever seen. They ended up being clearly highly motivated which is so important to me in teaching. When designing my jazz workshop and jazz improvisation course, I built it up in modules or layers which sit on top of teach other: meaning that you would start with something relatively simple and progressively build up your skills and knowledge. That way students are easily able to monitor their progress. I believe that anyone who plays an instrument is potentially able to improvise on it. But making progress is usually a matter of practising hard and dedicating the right amount of time to it. Hence the need for motivation.

Before embarking on jazz piano I started off with a classical education and that focused me on the mastery of the instrument. So I feel that being able to play your instrument well is a prerequisite for progress in playing jazz. I often started off with a great number of guitarists in my workshops. I would ask them to play scales, which they were not expecting, and this usually quickly led to a decrease in the number of guitarists in the class. They were expecting to play chords on the guitar but when it comes to soloing the guitar is just as much a horn as a saxophone and you need to be able to play lines. Typically, a lot of jazz piano includes playing chords in the left hand and tunes and improvised lines in the right hand. But the jazz pianist does not play chords as such but rather voiced chords, that is a selection and combination of notes which have developed as the music has evolved. Much of the left hand voiced chords have since been codified by Bill Evans.

The starting point of jazz for me is the Blues and it is at the heart of the music. From a technical point of view, it relies on the dominant seventh of classical music but this chord is used in a very different way. In classical music, the dominant seventh, as the name implies is a passing chord that gets you from the dominant to the tonic. However, its function in jazz is quite different. I prefer to call it simply the seventh and not the dominant seventh. The seventh is a chord which lies half way between the major seventh and the minor seventh each of which have their own matching scales. Similarly, the seventh has its own matching scale or scales which in some sense also lies between the "happy" sound of the major scale and the "sad" sound of the minor scale. The blues scale is simultaneous happy and sad.

After the blues, I introduce what I call the four routes to improvisation.

- The first route is called scalic. It uses the scale or scales which correspond to each chord, and they are used to create a musical line which is often connected stepwise, moving up and down the scale, starting and ending anywhere.
- The second route is chordal. Here you bounce up an down the different notes of a chord, in any order and in any pattern.
- The third route I call special devices. These are particular devices which work well on a piano, such as locking the hands together to play chords when improvising a line. However, they need to be used with discretion.
- The fourth route is called motivic. This is the key route. It usually involves taking elements of the tune and turning them into musical phrases. This is the hardest of the four approaches to employ. Here, the motif or melodic fragment you use is the key ingredient. Therefore, I usually start taking little elements of the tune in order to do something with them, such as translating them, inverting them, playing them backwards, or just making use of the rhythmic elements to shape a phrase. So one approach I use to develop this approach is something I call "hot licks". Here, I play a jazz phrase in time which the student is asked to repeat. I usually start with simple phrases and build up to playing the type of phrases that are typical in jazz. I make the distinction between "singing the play" where what you play should first be singable and you should also be able to sing along as you play and "playing the sing" where you compose the musical idea first and then you articulate it on your instrument.

As the course unfolds, I try to distinguish between what I would call "essential material" and "advanced material", I guess what teachers call "differentiation". The course is aimed principally at developing playing jazz piano with a rhythm section, as opposed to playing unaccompanied jazz piano. To this end I make use of "playalongs" - recorded backing tracks of bass and drums playing jazz standards. So, for example, the students can select a C blues and play along with the bass and drums, first playing a specific blues tune and then improvising on the chord structure. This is important because the fact of playing along with a recording forces the students to play in time - an essential element of jazz. It also forces the students to "keep up" with the unfolding chord sequence. The student also has the option to alter the speed and the pitch (key) of the playalong in practice.

In summary, using playalongs prepares the students to play in jam sessions or in groups with other musicians. Jam sessions have a special status in jazz. A jam session usually provides a rhythm section of bass and drums. Often there is a piano or guitar as well for horn players. Then the idea is that the guest musician plays together with the house rhythm section. The material used usually consists of blues and well-known jazz standards. The online jazz piano course covers over seventy such tunes. A student successfully completing the course should be happy either playing in a trio format or in a group format incorporating horn players. They should be able to play tunes in their own right, improvise over them and support other instrumentalists in doing the same. The joy of jazz piano playing should, I hope, be made available to them.

Acknowledgement

This text is based on an interview with Luc Steels, transcribed and edited by Maria Ferrer Bonet.

This page intentionally left blank

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-79

Chapter 6.
Learning To Be A Singer

Josep-Ramon Olivé I Soler
Guildhall School of Music and Drama, London

Abstract. What are the key challenges in becoming a professional singer of classical music? This chapter is written from the perspective of a professional singer who has been taught using the traditional highly personalised 'humanistic' approach to music education, which starts in childhood with choirs, and then focuses on advanced voice technique, development of musical memory, as well as the acquisition and application of linguistic and historical concepts in approaching a score. It raises the question is in how far on-line learning courses can be helpful to augment this kind of training and in how far ideas from a humanistic approach can be transported into the domain of on-line music learning.

Keywords. Voice training, humanistic education.

1. Introduction

I am an opera singer, currently continuing my lifetime ambition at the Guildhall School of Music and Drama with renowned professor Rudolf Piernay, in order to be able to become a professional opera singer and perform in the main opera theatres of the world. Having started a proficient singing career in my country and abroad has allowed me to be involved in many projects of different nature, such as recitals, opera productions, small ensembles and professional choirs.

I started my music education at the early age of 4, having a wide-ranged education involving many different aspects of music. I was enrolled in the "Escolania de Montserrat", a well-known children's choir located in the Abbey of the magnificent "Monsterrat" in Catalunya (Spain), which has a history going back up to the 13th Century. This experience allowed me to participate in several young singers' programmes, such as 'L'Académie Baroque Européenne d'Ambronay' or the 'Proyecto Pedagogico' in the Teatro Real of Madrid (Spain). I have also collaborated with numerous amateur and semi-professional choirs, which drove me into obtaining a bachelor's degree on choral conducting and singing, and after that towards a specialisation in singing with an MD.

Throughout history, there have been multiple attempts to define what elements make an opera singer. This chapter is not one more on the list, but rather a description of what experiences have influenced me to pursue a career as an opera singer, and which were the key factors that made it possible.

2. What a singer has to learn

Solfège is the most basic kind of musical education and the one I started with. It introduces you to pitch control and musical score notation, although it is a very wide subject and the difficulty increases with the years. The fact of studying solfège eventually allows you to sight-read, which is the ability of rapidly reading and performing a score, even without having read it before. There are a lot of teaching methods for solfège worldwide, although I believe that the fact of continuously learning new repertoire combined with solfège exercises is the most efficient way to reach strength and skilfulness.

Vocal technique is a requirement for any kind of singer, as it teaches how to properly use the organs related to singing, such as the tongue, vocal chords or the larynx. However, it is not only related to upper-organs, but it also includes the whole body, especially the thorax, which has a big impact on breath control. Breath control is one of the most discussed and controversial parts of vocal technique and there are many lines of study about how it should be used. Vocal technique allows singers to use their voice to the maximum and widen the range of styles, avoiding injuries and achieving an apparent ease at singing, even though it is extremely demanding. Vocal technique is applied differently to every singer, and every singer has a natural beauty, usually related to the spoken voice, which can't be modified.

Intonation, Articulation and Phrasing There's a blurry line between what is considered vocal technique and what is related to the musical aspect of singing, so some aspects of these properties are not clearly defined. Intonation is the capacity of singing in tune, the ability of adapting the pitch to each note. This ability, on the one hand, is related to the sound perception capacity, but on the other hand there might be technical problems that make it very hard for people to reach the right tune.

Articulation is influenced by a physical dimension and a musical dimension. Regarding the physical aspect, it is related to the position of the muscles and organs involved in pronunciation, so it is defined by the place an manner of articulation. Regarding the musical aspect, it implies how the musical note is sung, what's the attack and how long you hold the note (i.e. stacatto, legato, marcato...).

The concept of articulation often gets confused with phrasing, which is a concept related to the text and music, rather than the voice. Phrasing defines the melodic flow of the music piece, and it is a subjective concept, susceptible to discussion. As a singer, text phrasing must be defined first, in order to be coherent with the spoken language, then comes the musical phrasing. In the end, the phrasing will be a combination of both, chosen depending on the singer's criteria.

Language Diction and Text The main difference between singing and other disciplines lies in the text. The fact that singing uses two different languages, musical language and the language in which the text is written, intensifies the communicative effect. Therefore, the singer carries the responsibility of being precise and careful with the text's pronunciation. An opera singer is asked to use a wide variety of languages, such as Italian, French, German, English or Russian, and although the singer is not required to speak them perfectly, he should be able to achieve a native-like pronunciation through phonetics training. Note that diction when singing is not the same as in speaking, which also requires special training. The diction must be projecting and exaggerated, as well as comprehensible, as it must reach every corner of the room without any kind of amplifier or microphone.

3. Paths in musical education

Besides music education proper, there have been a lot of parallel music experiences that provided me with skills and knowledge that I now consider essential for my singing career. Here is an overview of those I believe had the largest influence.

Professional Opportunities vs. Education I have always believed there should be a stable balance between music education and professional opportunities. Certainly, a strong basis of knowledge and practise is required, but I recommend as well getting involved in non-academic projects as soon as possible. In my case, I signed in for as many projects as I could, regardless of time constraints and difficulty level. Clearly, nobody should do something they don't feel prepared for, yet in the end you learn from every project you take part in, sometimes what you should do, and sometimes what you shouldn't. Having the advice of teachers and other reliable people is also recommendable, as they have more experience and they can advise you of what is best at each stage of your career.

Early Music Early music is considered a section of classical music where music from the 15th to the 18th Century is performed. Nowadays, it is played following the contemporary musical practice treaties and using instruments of that period. These treaties intended to teach the performer how to play an instrument, how to ornament the music piece and how to reproduce the music faithfully and with the appropriate style. The fact of getting involved in early music has been very important for my singing career and it has allowed me to meet, among many other great musicians, one of the best early music performer worldwide, Jordi Savall. With him, I found myself constantly learning a new repertoire and performing with some of the best instrumentalists and singers, as well as inspiring me in many different aspects. For instance, Savall focuses a lot on the music's message, considering it of vital importance and encouraging the performers to deliver it confidently. Moreover, he insists that music is a channel that links people to what they are expecting from the music, which might be very different among the listeners (to entertain themselves, to connect with God, to achieve personal peace, etc), and this should be borne in mind by the performers.

Music in group is a concept that requires specific attitudes in order to succeed. My experience includes working with choirs and small ensembles, and there are four main lessons I learned from it.

1. First, cooperation. When performing in group, there are a lot of different people with very different perspectives and attitudes, so it is essential to find a common ground and a common intention in the aim of creating a performance together.
2. This leads to the second attribute, listening. In general, it is already an indispensable property of any singer, although when it comes to singing in group there must be a constant awareness of what the rest is doing.
3. In third place, precision is crucial to be able to put all these individual performances together and create only one.
4. Finally, adaptation is the last requirement for a successful group performance, as there should be a balance between all the members, keeping in mind that no member of the group is more important than the other, and all opinions should be taken into account.

Although this is general advice for group performances, there is a wide gap between a big and a small ensemble and each requires different skills and methodologies. In any case, small ensembles sometimes require even more attention and preparation than bigger ones.

Conducting a Choir Just like singing in group, conducting a choir requires special abilities that will determine the success of a performance. First of all, a conductor must transmit a sense of confidence to the group, an image of support, and this cannot be achieved without thoroughly preparing the piece. When it comes to preparing a performance, the conductor should be the most knowledgeable about the piece, as it is the conductor's responsibility to teach it to the choir.

However, achieving a correct performance of a piece does not only depend on the amount of preparation the conductor has had, but also on other skills that play a crucial role in its outcome. On the one hand, authority and leadership are necessary to be able to get the best out of them, not in a sense of severity and rigidness, but in order to engage the choir in the project and encouraging them to have the confidence mentioned earlier. On the other hand, listening is as important as in any other music field, but for a conductor it means understanding the key factors to unify the voices and create the characteristic sound of that particular ensemble. Moreover, a strong ability in listening will allow the conductor to identify mistakes and improve the individual performances.

Finally, having enough resources to make changes and knowing when and what changes must be applied to each case has an impact on the outcome of the performance and it can determine the success of it. In the end, however, in order to become a good conductor, practise is the best teacher and for me, conducting different-aged choirs was one of the most enriching experiences. Every choir has its specific need, but when there's an age difference the conductor's method and energy might considerably vary. For instance, conducting an amateur senior choir demanded patience and attentive explanations, as well as entertaining and joyful repertoire. Differently, I conducted an amateur young choir, which had a high music level, and it required me to be more demanding and to introduce challenging repertoire. Both experiences shared the feature of ”conducting a choir”, but, as it is shown, both required specific skills and techniques in order to succeed in a performance, which makes experience an essential part of an integral conductor.

Learning to Play Other Instruments Being a singer usually implies sharing the stage with other instruments or a full orchestra, hence why playing other instruments might be of the utmost importance, both for you and the instrumentalist(s). In my case, I learned to play cello and piano, and it proved to be a vital step for my singing career, especially piano, as it helps you to learn the music pieces and to understand the instrument's part of the accompaniment in the piece, therefore improving the merging process of both parts.

4. Humanistic education versus on-line learning

The “Escolania de Montserrat” is a boys' choir school located at the abbey “Santa Maria de Montserrat”, near Barcelona, Catalonia, Spain. The boys sing twice every day in the church, learning new repertoire every week and practising solfège on a daily basis. They learn music at a high level, quickly improving musical expression and sight-reading, by combining musical education with routine and enjoyment from the early stages of a singer. There's a focus on music formation, as the vocal formation will not be possible until the children obtain their ”mature voices”. Being a part of the “Escolania” has been the most important musical experience of my life and it has defined my career as a singer.

It seems to me unlikely that this type of humanistic education can be carried to on-line learning. On the other hand, the elite training is accessible only to a very select group that will become the top performers later. So the use of MOOCs, and particularly social MOOCs that foster cooperation within a community of learners, can still be a very satisfying way to broaden the opportunities for many.

Acknowledgement

This contribution is based on a talk presented at the PRAISE workshop in Casteldefells (Spain), december 2014. It was described and edited by Maria Ferrer Bonnet.

This page intentionally left blank

Part III

Platforms for Learning Communities

This page intentionally left blank

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-87

Chapter 7. Music Circle: Designing Educational Social Machines for Effective Feedback

Matthew Yee-King, Maria Krivenski, Harry Brenton, Mark d'Inverno
Goldsmiths, University of London

Abstract. We report on our development of an educational social machine based on the concept that feedback in communities is an effective means to support the development of communities of learning and practice. Key challenges faced by this work are how best to support educational and social interactions, how to deliver personalised tuition, and how to enable effective feedback, all in a way which is potentially scalable to thousands of users. A case study is described involving one to one and group music lessons in an on-campus, face to face, higher education context that were observed and analysed in terms of the actions carried out by the participants. The actions are described and it is shown how they can be formalised into a flowchart which represents the social interactions and activities within a lesson. Through this analysis, specific scenarios emerged where the feedback being given might not be effective, e.g the recipient not understanding the feedback or the provision of feedback which is not specific enough. In answer to these scenarios of ineffective feedback, the requirements for a technological intervention which aims to make the feedback more effective are proposed. With this in mind, we are then able to describe a novel technological platform which has been developed as part of a large-scale European research project and which aims to support effective feedback. The platform is based around focused discussion of time based media, embedded within existing teaching activities at a research led higher education institution in the UK. We outline how it is being used in a blended learning model to support the teaching and learning of music. We reflect on the experience of developing techniques and systems for enabling communities of e-learning and describe our evaluation methodology which involves several case studies and approximately 400 users in its current phase.

Keywords. Music circle, collaborative learning, MOOCs, online learning, feedback

1. Introduction

Our research project is concerned with the development of a *social machine* which aims to support and enhance the experience of learning music through the optimal provision of feedback. Key challenges we face in this work are how to support educational and social interactions, how to deliver personalised music tuition, and how to enable effective feedback, all in a way which is potentially scalable to thousands of users.

In this paper, we present our method for addressing these challenges through an initial period of teaching observation and analysis followed by the development of a technological platform via a participatory design process. The methodology is summarised in Figure 1. Following that, two key research outputs are presented: an analysis of one to one and group music tuition within our institution and a novel e-learning platform we have developed in response to this analysis. The teaching analysis resulted in a list of archetypical teaching and learning activities, shown in table 1, an ontology of musical feedback, shown in Figure 3 and flowcharts describing interactions within lessons as shown in Figure 2. The technological platform is essentially a repository for audio and video recordings which allows the user to upload media then to share it with communities of other users who can then place comments relating to the media along a timeline. It is described as a set of system requirements in table 2 and as screenshots of its media discussion interface and social timeline in Figures 4 and 5.

1.1. Background

Let us first consider what we mean by a social machine. Tim Berners Lee is credited with having coined the term social machine in 2000:

> Computers can help if we use them to create abstract social machines on the Web: processes in which the people do the creative work and the machine does the administration [1].

This quote is contextualised in the transition to web 2.0 where the process of publishing content and interacting online was democratised with technologies such as blogs, social networks and so on. In 2013, we find ourselves in the age of the social machine, where the point of interest for internet technologies is no longer the architectural underpinnings but the way in which people and machines interact within these systems. De Roure et al., are concerned with the observation of these social machines and provide some examples: Wikipedia, Ushahidi, Galaxy Zoo, reCAPTCHA and Mechanical Turk [16]. Moving to the educational context, 2012 was the ‘year of the MOOC’ [14]; indeed, one of the authors of this paper ran a MOOC with an enrolled student body of 97,000. With their extreme student to staff ratios, MOOCs rely upon interactions between peers for support and assessment; this is a level of social interaction that seems beyond what has been seen previously within standard VLEs. Since they are technological systems supporting a range of social interactions, we consider them to be another example of a social machine.

Now let us consider the term ‘feedback’. We define feedback in the educational context simply as a *reaction to a learner’s output which is somehow made visible to the learner*. In higher education in general, feedback is considered very important. It is one of the key areas covered by the UK National Student Survey and historically one of the lower scoring areas in terms of student satisfaction [7]. So feedback is important and is not always being done well, but how can we do it better? Juwah et al. present a list of 7 principles of good feedback in higher education, wherein good feedback

- Facilitates assessment (reflection) in learning
- Encourages teacher and peer dialogue around learning
- Helps clarify what good performance is (goals, criteria, expected standards)
- Provides opportunities to close the gap between current and desired performance,
- Delivers high quality information to students about their learning
- Encourages positive motivational beliefs and self-esteem and

- Provides information to teachers that can be used to help shape the teaching

[8].

These are useful general principles but music education is a specific case where the contexts and nature of feedback are perhaps quite different. Therefore, in this paper we will present our analysis of feedback within music education with specific examples, then show how we have developed a technological system which aims to support that specialised kind of feedback.

1.2. Previous work

In this section, we will provide a brief overview of some related work in the areas of social discussion of media, online music education and peer interactions. The platform provides a media repository and timeline based discussion functionality; a similar commercial platform is Soundcloud, which allows users to maintain and share a repository of audio files and to post comments to a timeline [2]. Considering the concept of annotations placed on a timeline, Latulipe discusses various projects using timeline based discussion systems including the 'Video Collaboratory' [9]. Puig et al. developed the 'Lignes de Temps' software which provides a multitrack timeline aiming to promote polemical discussion [15]. Moving to the music education area, there are a range of commercial online platforms such as ArtistWorks [10] and Berkley Online from the Berklee School of Music. Indeed Berklee have been running musical MOOCs on the coursera platform, using SoundCloud for peer discussion [12]. There has also been significant public research undertaken into technology for music education, such as the European funded i-maestro and VEMUS projects, both of which focused in part on the specificity of feedback [13], [6]. The concept of social interactions between students within VLEs did not arrive with the xMOOC in 2012, of course; the cMOOC which came before it had perhaps a more radical, distributed pedagogy [18]. Going further back, forums have been a standard component in VLEs for a long time and new types of VLEs emphasising social interactions have been reported in the literature. For example, Shi et al. describe their Topolor system which enables 'social personalized adaptive e-learning' [17]. Finally, to contextualise our methodology, we use a grounded theory approach to analyse our lesson observations and a participatory design approach to develop the features of the platform [4], [11].

1.3. Research Questions

Our research project has several high level research questions:

1. How well does our approach increase participation in musical learning activity?
2. How important is giving and receiving feedback online for engagement with practice?
3. How do we correlate engagement and feedback in a community?
4. What is the right level of social coordination and structure that students want for online-supported learning? Can we provide interfaces for non-technical people to design social coordination?
5. How can we evidence musical competencies and musical development in students?

6. How can automatic techniques be used to evidence feedback in music learning?

In the work presented here, we describe our 'approach' and provide evidence about the nature and importance of feedback which underpins several of the questions above. We also provide answers to how one might evidence musical competencies.

1.4. Structure of this paper

The background and motivation for the work has been presented in this section. In section 2 we will describe our methodology for building social machines combining teaching observation and participatory design. In section 3 we present the outputs of the methodology including the observed teaching and learning activities, types of feedback and a description of the features of our new platform. In section 4 we describe the ongoing evaluation of the platform with 400 users. The paper ends with a discussion and conclusion in section 5.

2. A Methodology for building social machines combining teaching observation and participatory design

The development of our platform has taken place in 4 phases. In phase 1, *teaching observation*, we observed and recorded 23 undergraduate instrumental and vocal performance lessons at our institution. The lessons involved 9 teachers teaching guitar, voice, piano and group and 14 individual students. The lessons were in either one to one or group format and spanned the popular and classical music degrees. Recordings of the lessons were transcribed to approximately 500 pages of text and notes were taken by the researcher observing the lessons. In phase 2, *analysis*, a grounded theory approach was used to code the activities within the lessons in order to identify key teaching and learning activities. This approach 'fosters seeing your data in fresh ways and exploring your ideas about the data through early analytic writing' [4]. The activities were then organised into higher level descriptions in the form of flowcharts describing different types of lessons. A particular emphasis was placed on the flow of feedback between participants in these lesson archetypes. In phase 3, *basic requirements*, we drew up some basic requirements for the platform in order for it to support the teaching effectively. This would allow us to bootstrap the basic functionality of the platform ready for the next phase. In that phase, *participatory design*, we used a participatory design approach, where the input of users is sought and acted upon throughout the iterated development lifecycle [11]. In a sense, the final phase includes its own observation, analysis and requirements phases, except that the observations are of users using the system (for real teaching and learning). This final phase is ongoing. Figure 1 illustrates the relationship between the 4 phases.

3. Outputs from the Methodology

In this section we will present the outputs generated by the teaching analysis and participatory design process.

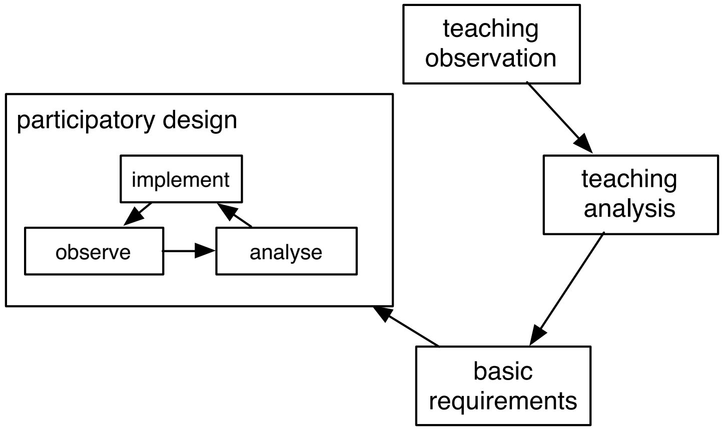

Figure 1. The 4 phases of platform development

3.1. Enumerating teaching and learning activities

We were able to identify 9 distinct teaching and learning activities from our lesson transcripts and observation notes and these are shown in table 1 with examples of each from the transcripts.

3.2. Teaching workflows

Through our lesson analysis, we were able to identify lesson archetypes which appeared several times in the observations. We call these archetypes 'teaching patterns', after Eckstein et al. [5]. A complete description of the teaching patterns is beyond the scope of this paper but a single example flowchart representing a lesson where a student performs in front of their tutor and peers can be seen in Figure 2.

3.3. Feeding back about music

Perhaps unlike some other subjects, there is a rather discrete and finite ontology underlying the types of feedback one might receive about playing a musical instrument. As part of our analysis, and based on previous work, we have developed a detailed ontology to describe feedback on musical performance, shown in Figure 3. It should be noted that we have identified two broad types of feedback: firstly, feedback connected to desirable traits in a musical performance, as shown in the majority of Figure 3 and secondly, 'information for guiding tactics and strategies that process the domain specific information' after Butler and Winne [3]. The latter might also be expressed as encouraging the learner to develop their self reflective skills, their *inner teacher*.

3.4. An understanding of problems with feedback provision motivating essential platform requirements

We now have a clear idea of the context within which feedback is given (e.g. lesson flowchart in Figure 2) and the expected content of that feedback (i.e. the ontology in

Table 1. The 9 distinct teaching and learning activities in one to one and group music lessons

Activity	Description	Example
Transmission	Tutors provide theoretical and practical information to students	so whatever you do to your mouth, it's the same sound because the tongue is going right up against the soft palette, so the sound can only come out in your nose.
Performance modelling	Tutors or students perform good and bad examples of extracts from a composition	A musical activity
Identify and solve	Identify, discuss then suggest solutions to performance problems.	Okay, did you hear that? The music is very uneven... Let's experiment a bit. Let's do it this way. I'll play the right hand with you the first time. I am going to go for just a legato version. Then you will have a go at it hands together and I would like you try to a legato version so then you are not affected by the separation of the notes.
Practicing solutions	Students put the solutions from the identify and solve activity into practice in their playing	A musical activity
Feeding back	Self, peer and tutor feedback on a performance, after it has happened	That's fine, that sounded pretty good. The very first time it sounded - your down beat sounded a little bit like 'oh this is a down beat, I'm going to play loud now.' Always be careful about how you're shaping it.
Checking student understanding	Initiated by student or tutor, student understanding is verified through dialogue	[Tutor] From there, just flatten the 3 and you've got Dorian and add to that flatten the 6, you've got Aeolian, if you want to continue, what would you do next? Anybody know? [Student] Flatten the 2? [Tutor] Exactly right! Flatten the second, becomes? [Student] Phrygian. [Tutor] Phrygian, that's right! Which is a very nice scale, I'm fond of it.
Discussion of goals and ideas	Discussion and negotiation of assessment or other goals and creative ideas	[Teacher] What is romantic for you? Let's engage in this kind of discussion. What is romantic? It's important. What is romantic for you? [Students]: To express your emotions, along with that establishing a connection. [Tutor] Don't you think that being romantic also sometimes can mean trying to be a bit more individual than you normally are in the real world, to be more special?
Performing	Students performing a prepared piece	A musical activity
Directing	Tutors verbally guide a student performance in real time	[Teacher] Top string this time. Take that off so you're playing - you want that note. There's G. Put your little finger back. G7. Put your finger back. The difference where your first finger is, yes, that's suspended, that's G. You can hear it.

Figure 3). However, we were able to identify several reasons why feedback might not be effective, listed below. Note that at this point, we begin to consider the basic requirements for our platform which will allow it to address these problems directly.

1. The underlying ontology driving the feedback is not well understood
   The platform should be able to gradually expose an ontology in a range of ways. (e.g. through suggestion of relevant terms, and the provision of automated, high level annotations)

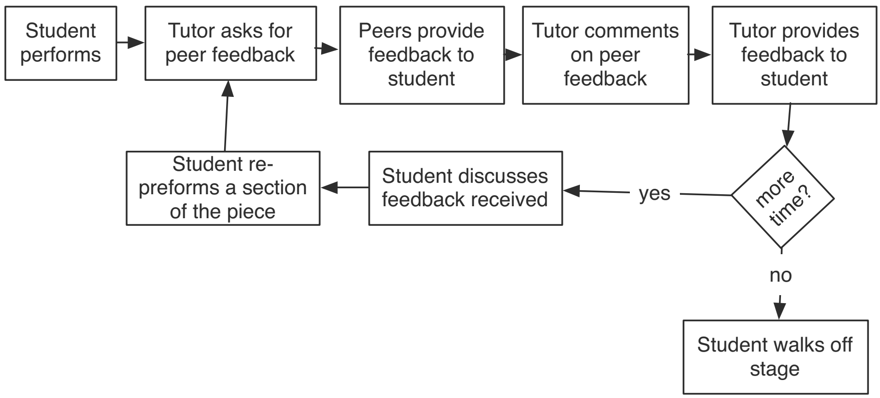

Figure 2. A flowchart describing a peer feedback lesson where a student performs in front of their peers and tutor, then receives feedback.

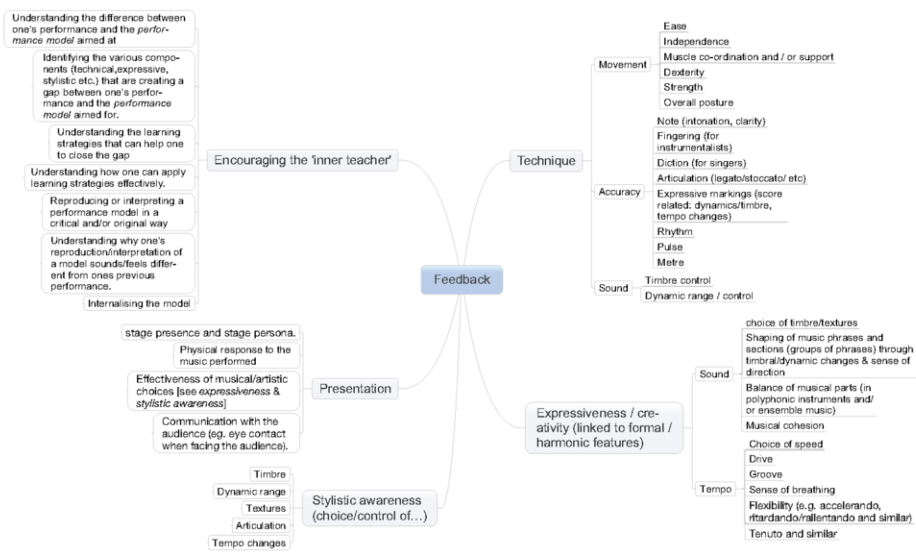

Figure 3. The ontology we developed for describing feedback about musical instrument playing.

2. The feedback is not remembered
   The platform should make feedback easily accessible for later reflection, not hidden away in a forum somewhere, for example.
3. The tutor is the sole source of trusted feedback
   The platform should embody a community of learners pedagogy, to emphasise the value of feedback from peers and tutors alike.
4. The feedback given to peers is not honest, e.g. 'too nice'
   By building a platform that enables more precise feedback related to a specific

ontology, feedback should naturally become more honest, as the emphasis for the feedback is aimed away from the individual and towards particular aspects of a performance.

5. The relevance of feedback to a particular performance is not understood
   The platform should encourage the provision of feedback which is specific and well justified.
6. The feedback is too narrow
   Here, the feedback focuses on a limited part of the ontology, typically due to time constraints in a lesson. The platform should encourage a community discussion around a greater number of performance aspects.

3.5. Platform design

The final phase of our methodology was the iterated development of the platform. This process is ongoing, but it moved through 8 versions during the first year, where increasing numbers of users were involved at each stage. The resulting platform is essentially a repository for audio and video recordings which allows the user to upload media then to share it with communities of other users who can then leave comments along a timeline. Its key features are listed in table 2 and shown in figures 4 and 5. At the end of this first year of development, the system was in active use within 5 undergraduate modules at 2 institutions. In the following passage, the key features and motivations for their inclusion will be discussed.

Easy access, personal media repository. The aim is to remove barriers to content uploading and sharing and to make content easily accessible for later review. The platform includes simple record and upload apps for iOS and Android to make content addition as easy as possible as we identified that the often over-complex process of putting content into VLEs can be a serious barrier to uptake for students.

Simple sharing and community model. The aim is to increase user confidence in uploading and sharing media. The platform provides a very clear method of controlling who the content is shared with. Also, users can delete any comments made about their content.

Intuitive discussion interface with content prompting. This feature aims to motivate commenting activity and to encourage use and understanding of appropriate terms from the ontology.

Social timeline. Feedback is always connected to a particular range of time in the media. Also, all commenting users have individual timelines displayed below the media. This promotes awareness of the community opinions, making feedback specific to a person and a time

Powerful discussion system. Users can reply with audio, video, text and so on. Audio and video responses within the platform can then become a subject for discussion in themselves, with their own social timeline.

Automatic feedback agents. We are developing software agents which are able to feedback automatically about musical performances. They work by comparing different performances and making high level comments about the variations, connected to the feedback ontology. This provides a ‘neutral’ source of feedback and exposes the learner to the ontology. A full description of the feedback agent is beyond the scope of this paper but it is built around machine learning and audio analysis techniques.

The key features of the platform are listed in table 2, where we also compare them to the closest equivalent commercial system, SoundCloud and a well known open source e-learning tool with social features, Mahara.

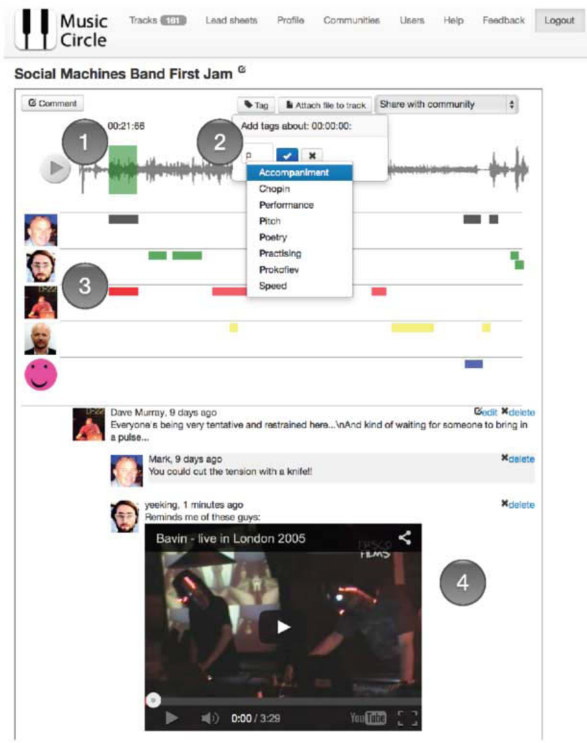

Figure 4. The music circle media discussion interface. 1) The waveform display, showing a highlighted region, 2) The tagging dialogue, showing a drop down list of pre-used tags 3) The social timeline, showing sets of time linked comments created by several users 4) a discussion thread based on a single region in the recording, including an embedded youtube video.

4. A description of our ongoing evaluation with 400 users across 2 institutions

The participatory design process aims to suggest then optimise platform features. In a sense, this represents an ongoing, evaluation and improvement cycle. However, as stated in the introduction we are interested in the evaluation of social machines and the activities they enable at a higher level than basic platform features. In this regard, we are running significant case studies with our platform with approximately 400 users spread across 2 institutions and 5 different modules. The evaluation scheme consists of qualitative and quantitative methods. In particular, we will be using interviews, survey tools and user activity metrics including social network analysis. This will allow us to address the research questions listed in section 1.3 with a variety of perspectives. We anticipate being able to analyse a data set containing hundreds of media items, thousands of comments and many thousands of interactions.

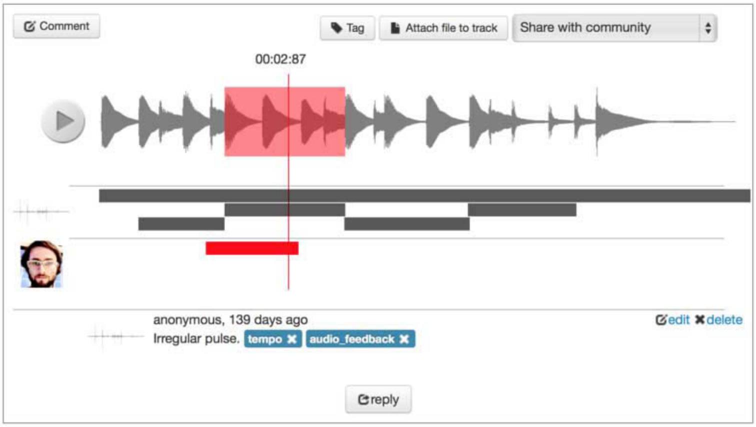

Figure 5. The social timeline, showing sets of annotations from two users. Each block in the timelines represents an annotation connected to a specific region in the recording. Here, the top timeline was created automatically by the feedback agent.

Table 2. The key features of the system compared to some pre-existing systems which we have used for teaching at our institution

Feature	MusicCircle	SoundCloud	Mahara
Easy access, personal media repository with mobile media capture client	x	x	
Simple, transparent sharing and community model	x		x
Intuitive discussion interface with content prompting	x	x	
Social timeline with region selection	x		
Powerful discussion system	x		x
Automatic feedback agent	x		
Suitability for use as a research platform (data access, privacy etc.)	x		x

5. Conclusion

Work has been presented which faces the challenges of how to support educational, social interactions, how to deliver personalised music tuition, and how to enable effective feedback. A methodology for addressing the challenges has been described which takes real observational data and analyses it into formalisations of teaching and learning activities. The outputs from this methodology have been presented, including a list of key teaching and learning activities, a flowchart describing the interactions within a typical lesson, and an ontology of types of feedback. It has been shown how the outputs have been iteratively interpreted into the design for a novel e-learning platform driven by social interactions and effective feedback. The current system has been introduced and the ongoing evaluation with 400 users has been described. The immediate targets for our future work are to increase the number of learners operating within the platform, to conduct an investigation of the wider applicability of the system, for example as a means to deliver recordings of lectures and the development of our tool kit for quantitative evaluation of the system. Inspired by the examination of the importance of feedback presented here, the longer term goal is to develop a deeper understanding of the nature and importance of feedback in the learning and creative processes.

5.1. User Evaluations

Table 3. Case study evaluations of the Social Timeline

Case study	Users	Evaluation
Undergraduate music students at Goldsmiths	49 first-year music students	User diaries, ontology of user comments
LCO (London Contemporary Orchestra)	X school children aged 13-15	Interviews with LCO staff
Jazz in performance at Leeds	14 undergraduate Jazz students	Reflective account from the teacher

5.1.1. Using the social timeline to support different user needs

The social timeline can encourage different types of desirable behaviour. For example, the teacher in case study 1 wanted her students to give an evidenced-based critique of their peer's work and avoid value judgements such as ‘I like this bit’ or ‘nice one. The Social Timeline helped achieve this by requiring users to link comments to a selected region of audio. This constraint forced the undergraduates to justify and substantiate their opinions with evidence. For example, one student made the following comment.

> ‘Yeh I agree [name of student] - I’m holding the pedal down way too long through bars. Looking back I’m also noticing that the dynamics could be slightly more exaggerated and the top note of the chords (which follows vocal melody) could be brought out more”(Undergraduate 2).

To which the other student replied.

> ‘Accompanist, be careful with holding that sustain pedal for too long because it can create a muddy feeling and this part should be quite clear to reflect the lyrics and tone of the vocalist’s voice”(Undergraduate 1).

Giving and receiving this type of evidence-based feedback was initially challenging and daunting. But over time it helped to build group cohesion and gave insight into peers previously hidden creative processes. In contrast the teacher in case study 2 actively encouraged value judgements such ‘I like this’ and ‘I thought this was good’ and did not want users to give individual critiques.

> ‘I think so far in our project, we have been working in very broad strokes. We’ve been working in getting people playing together and in feeling comfortable, and generating lots of material. A lot of our work in this project is about positive affirmation and about taking the ideas that are given. Even if comments do become more specific as the rehearsals progress, they are still likely to be directed at the level of the group rather than the individual’(Undergraduate 1).

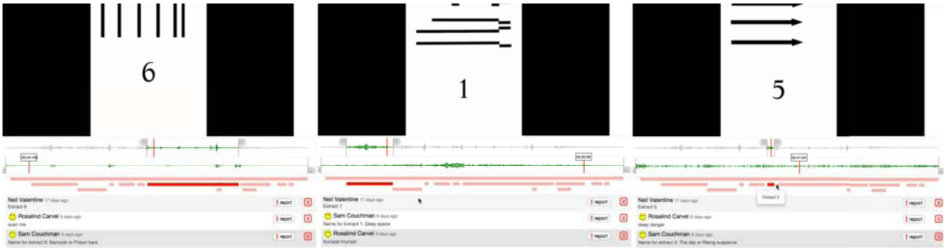

Figure 6. Name the shape challenge

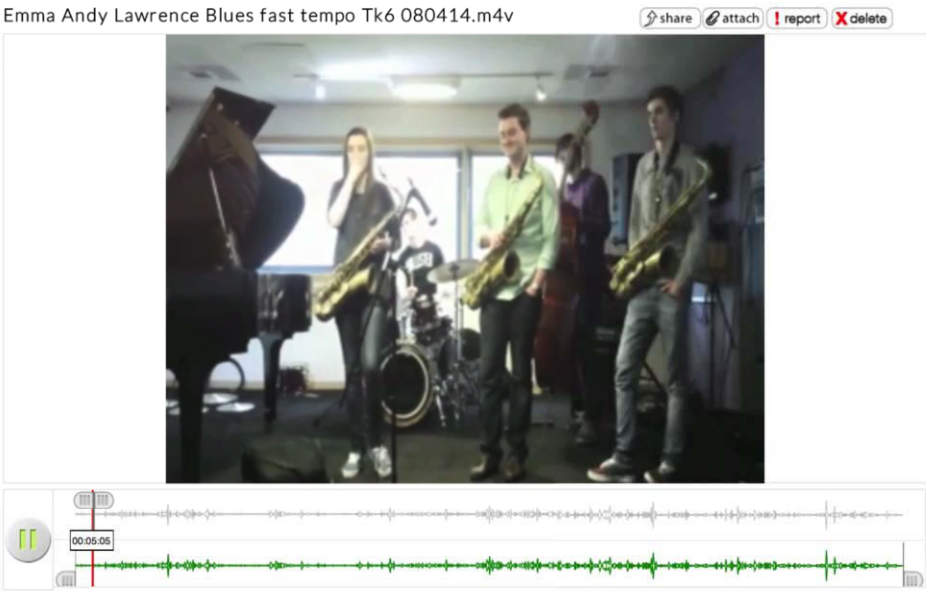

Figure 7. Jazz in performance case study

To achieve this, the teacher set challenges such as 'name the sound-shape' which directed feedback towards an external object and away from peers' performances.

The Jazz teacher in case study 3 wanted his students to understand the integration between theory and practice (see chapter X in this book).

> 'I strongly believe that one of the problem areas in Jazz Education is that the dissemination of theoretical knowledge is very often given precedence over, or divorced from, the practical application of theory to practice (the doing). My aim in the module was to create this kind of environment where a multitude of performance,ensemble, conceptual, theoretical and technical issues could be exposed and explored within a working rehearsal/Jam session situation.'(Jazz teacher).

To achieve this, he framed group discussions around videos of jam sessions.

> 'Video is an excellent medium for highlighting several layers of information in regard to performance. in addition to documenting the audio of a live performance, video captures complex interactions of information which flow and fold in on each

other to produce new and sometimes unexpected improvised dialogues and responses...MusicCircle accommodates these multiplicities and allows a variety of information strata to be disseminated, discussed and elaborated upon in ways that conventional classroom activity cannot possibly deal effectively within real time. The timeline allows both student and teacher to return to specific event areas and go into them at depth. Thus in turn creating more dimensions to the learning experience. '(Jazz teacher).

These three examples demonstrate the flexibility of the social timeline. It concentrates upon one central interactions : group annotation, and then allows teachers to structure teaching as they wish.

5.1.2. Conclusions from user studies

We drew five main conclusions from the user studies.

1. The social timeline encourages users to associate feedback directly with the creative object (music recording, video or visualisation). This helps to focuses feedback directly on the 'performance' rather the 'performer'.
2. The social timeline uses spatial constraints in the user interface that force users to justify comments with evidence, leading to the perception that feedback is more unbiased' and objective.
3. Videos capture complex interactions of information which flow and fold in on each other to produce new and sometimes unexpected improvised dialogues and responses from students
4. The social timeline makes working processes explicit and exposes them to public view. Seeing this previously hidden information can help students to overcome anxieties about sharing their own work.
5. It is very hard to build an active social network from a 'start'. Our experience led to the following recommendations: a) Assess contributions, if this is appropriate for the teaching context; b) Try to anticipate and remove organisation from day one (eg. getting users logins); c) Showcase curated content on the front page to promote good material; d) Embed Music Circle into students and teachers workflows and working practices.

References

[1] Tim Berners-Lee, Mark Fischetti, and Michael L Foreword By-Dertouzos. *Weaving the Web: The original design and ultimate destiny of the World Wide Web by its inventor*. HarperInformation, 2000.

[2] Andrew Bird. soundcloud.com, 2014.

[3] Deborah Butler and Philip Winne. Feedback and self-regulated learning: A theoretical synthesis. *Review of educational research*, 65(3):245–281, 1995.

[4] Kathy Charmaz. *Constructing grounded theory: A practical guide through qualitative analysis*. Sage Publications Limited, 2006.

[5] J Eckstein and Joseph Bergin. Patterns for active learning. *Proceedings of PloP*, 2002.

[6] D. Fober, S. Letz, and Y. Orlarey. VEMUS-Feedback and groupware technologies for music instrument learning. 2007.

[7] HEFCE. National Student Survey Findings and trends 2006 to 2010. Technical Report April, Higher Education Funding Council of England, 2011.

[8] Charles Juwah, D. Macfarlane-Dick, Bob Matthew, David Nicol, David Ross, and Brenda Smith. *Enhancing student learning through effective formative feedback.* Number 68. The Higher Education Academy, 2004.

[9] Celine Latulipe. The value of research in creativity and the arts. In *Proceedings of the 9th ACM Conference on Creativity & Cognition - C&C '13*, page 1, New York, New York, USA, 2013. ACM Press.

[10] Mike Marshall, Bryan Sutton, Martin Taylor, Jason Vieaux, Nathan East, Missy Raines, Ricardo Morales, D J Skratch, and David Bilger. Artistworks.com, 2014.

[11] Michael J Muller and Sarah Kuhn. Participatory design. *Communications of the ACM*, 36(6):24–28, 1993.

[12] Carin Nuernberg and Alex Perrier. Behind the Scenes with MOOCs: Berklee College of Musics Experience Developing, Running, and Evaluating Courses through Coursera. Technical report, 2013.

[13] B. Ong, K. Ng, N. Mitolo, and P. Nesi. i-Maestro: Interactive multimedia environments for music education. 2006.

[14] Laura Pappano. The year of the MOOC. *The New York Times*, 2(12):2012, 2012.

[15] Vincent Puig and Alexandre Monnin. a collaborative polemic-based video annotation platform. Technical report, 2006.

[16] David De Roure, Clare Hooper, Megan Meredith-lobay, Keble Road, Oxford Ox, Kevin Page, Don Cruickshank, and Catherine De Roure. Observing Social Machines Part 1 : What to Observe ? pages 5–8, 2013.

[17] Lei Shi, Dana Al Qudah, and Alexandra I Cristea. SOCIAL E-LEARNING IN TOPOLOR : A CASE STUDY. In *IADIS International Conference e-Learning 2013*, pages 57–64, 2013.

[18] Becky Smith and Min Eng. MOOCs: A Learning Journey. In *Hybrid Learning and Continuing Education*, pages 244–255. Springer, 2013.

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-101

Chapter 8. Giant Steps in Jazz Practice with the Social Virtual Band

[a] *Sony CSL Paris, 6 rue Amyot, 75005 Paris, France*

Mathieu Ramona [a], François Pachet [a], Stanislaw Gorlow [a]

Abstract. This chapter deals with the issue of learning how to improvize. Traditional MOOCs provide jazz students with comprehensive theoretical and motivate students to practice intensively on their own. However, without a view of one's progress, and without feedback, individual practice is a long and winding road along which many students get lost. Indeed, most jazz learning systems lack these two crucial ingredients, resulting in high drop-out rates. This chapter addresses the issue of designing tools for supporting improvisation practice by bringing in a social dimension enabling peer-to-peer feedback, as well as a cloud-based infrastructure enabling arbitrary vizualisations of the evolution of the student performance. We introduce *Social Virtual Band*, a system that lets learners improvize solos on dynamically created accompaniments, and that records and archives all the training sessions along with the provided accompaniments on the cloud. Simple automatic feedback is presented to measure the evolution of his skills, based on a comparison between played note with scales obtained from automatic harmonic analysis. We describe the overall infrastructure underlying such a tool and discuss how such infrastructure opens up new possibilities for learning music.

Keywords. Jazz, Practice, Cloud, Automatic music accompaniment

Introduction

Jazz improvisation is a skill that requires both formal musical training, notably in harmony theory and melody development, and a lot of practice time (the 10,000 hour rule [9]) to integrate and literally embody the knowledge. Embodiment is crucial in jazz improvisation because the requirement of producing and playing music in real-time make it impossible to think at the symbolic level. The skills necessary to build solos, and to *play the changes* are unique in that respect, and have even been considered by some authors as key for a wide range of activities, including managerial ones [12]. Practicing jazz consists essentially in playing solos over existing jazz standards. The acquisition of solo building skills is difficult to formalize precisely but is acknowledged to require both an intimate mastery of the instrument [27] and the ability to communicate musically with others ([7]).

Several solutions have been proposed to assist jazz improvisation practice. The oldest one is probably the minus-one recordings, notably the Aebersold series [13]. With minus-ones, the student plays a recording of the accompaniment of a jazz standard, and can freely improvise on it, usually with top quality backing. Though immensely successful, this solution can be used only for a predefined set of recordings, and cannot be costumized either tempo-wise, style-wise or structure-wise (i.e. number of solos). Sotware solutions have been developed, such as the famous *Band-in-a-box* system [15], which generates medium quality accompaniments of arbitrary tunes, in various styles. More recently, apps running on portable devices have been developed, such as *iReal Pro* [28], and also provide medium quality accompaniments on arbitrary chord sequences. In all these cases however, the performance produced is lost, no feedback is given to the musician so there is no possibility of reflecting on past performances, comparing them, or building a synthetic view of the student's skills evolution.

This chapter describes a system that attempts precisely to push jazz practice a step beyond conventional practice, by providing musicians with a cloud-based system that enables not only the collection and analysis of performances, but also the production of feedbacks, either automatically or by peers, which are essential to guide the learning process.

1. Improvisation Practice on Jazz Standards

1.1. Structure of a lead sheet

The lead sheet is the basic element that defines a song in jazz. A lead sheet is a combination of a melody (usually monophonic) with a sequence of chord labels defining the harmonic progression. This harmonic progression can be seen as a realization of an underlying functional path which gives a particular flavor to the melody. The lead sheet also includes a time signature (e.g., $3/4$ or $4/4$) and possibly stylistic indications that give information about the tempo and the rhythmic patterns to be used. Figure 1 shows an example of such a lead sheet.

A trained musician can infer a lot of information from the lead sheet: the global tonality of the tune, the modulations (through the analysis of the chord sequence), the harmonic function of each chord in the progression, and consequently the scales which fit with the chord sequence that can be used for improvising. A lead sheet represents the *essence* of the song, and constitutes a reference knowledge shared by all jazz musicians ([7]).

A typical jazz player generally starts by playing the melody (aka the *theme*) and then elaborates his solo by paraphrasing the melody, while respecting the harmonic and structural constraints. In traditional jazz (typically be-bop) ensembles, the drum, bass and piano handle the rhythmic part and provide the harmonic support, while wind instruments play the solo parts in turns. Later extensions of the traditional jazz canvas, such as modal jazz, have progressively loosened up the traditional canvas and now any instrument can play solos in turn, including drums and bass.

1.2. The LSDB database

The first requirement of an online jazz practice system is to access digital representation of lead sheets. We use the comprehensive online database of jazz lead sheets [21] which contains all published jazz lead sheets, and is the first large scale database with both
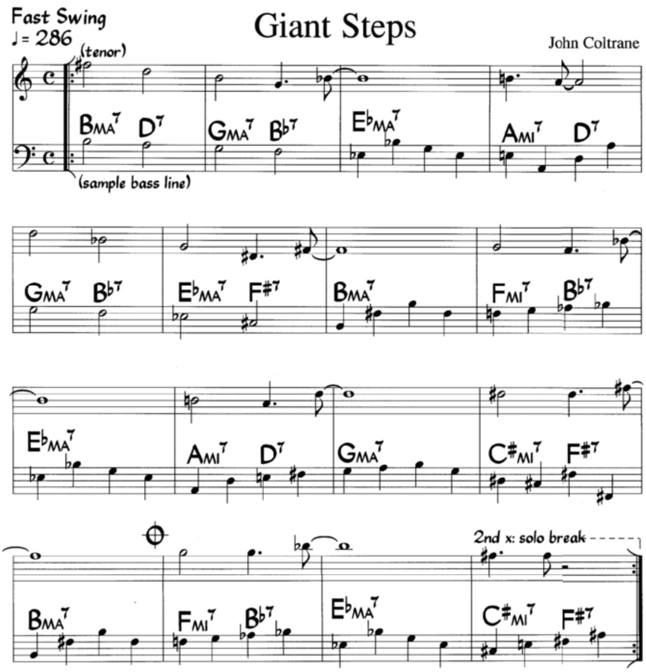
Figure 1. A lead sheet for jazz standard *Giant Steps* by John Coltrane.

melodies and chord progressions in electronic format. The database currently contains about 11,000 songs coming from 62 song books, of which 20 are fake books or compilations of jazz standards, 32 are Bossa Nova song books, and 8 are composer specific song books (such as John Coltrane, Thelonious Monk or Bill Evans).

Such a database is a precious resource for Music Information Retrieval [3] and automatic music analysis [8] [11]. In the context of this experiment, it provides us with a comprehensive set of lead sheets for generating accompaniments to support the practice of jazz improvisation over nearly any standard a musician could choose to play on.

1.3. Features of a solo

The improvised solo is classically a paraphrase of the melody, that can lead to multiple sorts of variations (transposition to different keys or different scale, slight rhythmic variations, repetition over different chords) on motives (defined as the smallest melodic entities) possibly chosen from the melody or even cited from other sources.

What makes a good solo is of course a complex question, because it can be evaluated according to many musical dimensions. Jerry Coker states [4] that five factors concur in the outcome of a remarkable improvisation: intuition, intellect, emotion, sense of pitch and playing habits. Learning improvisation consists basically in developing a conscious control, through the intellect, of the other four factors, which should become inconscious with practice. From the point of view of the listener, a good improvisation can also be defined as a very subtle balance between predictability, as a mean to create attention, and surprise over the listener's expectations, to avoid boredom.

In fine, automatic feedback necessarily relies on *features* that can be extracted from the audio signal by pure computation. The extraction of so-called *audio features* (i.e. numeric or symbolic values that can be deterministically calculated by a machine from an audio signal) is indeed a key issue in Music Information Retrieval [22], but low-level features (estimated with few computation steps from the signal, e.g. spectral centroid, temporal moments, RMS, etc.) are generally unable to catch perceptively relevant features. Although some approaches [19] can automatically infer high-level features that fulfill a given task, we focus here on the musical features that can be defined from a reliable transcription of the solo (i.e. the automatic extraction of the notes).

Supposing that the transcription is reliable, each played note is detected and characterized by its *pitch* and its *start* and *end times*. The following features are easily estimated:

- **In-scale rate**: the so-called *In-scale rate* measures the proportion of notes played in the scales expected from a harmonic analysis of the chord progression (this will be developed in Section 2.5). Figure 2 shows two examples of simple melodic phrases played on a D minor chord ; example (a) is perfectly in scale, while (b) has 33 % notes off scale.

(a) In-scale rate: 100% (all notes are in scale).

(b) In-scale rate: 66% (3 notes out of 9 are off scale).

Figure 2. Two examples of in-scale measure over simple melodic phrases played on a D minor chord.

  Of course a 100 % rate is not necessarily an objective since *playing out* (as it is often called in jazz jargon) is precisely a key part of what makes a solo enjoyable, for example through side-slips (see [14] for a pedagogical definition). The estimation of the expected scale in also a non-trivial issue [26], although as a first approximation, an ad-hoc scale can be associated with each chord of the sequence.

- **On-beat rate**: measures the mean time shift of notes close to the first and third beats. Of course, discriminating accidental time-shifts (out-of-beat notes) and intentional syncopation is a key issue for such a measure.

- **Continuity**: melodic continuity is considered a key aspect of a good *sense of melody*. It is a difficult challenge for a human when playing fast, as it requires the ability to find quickly short paths between the note currently being played and the next ones, which may be in a different scale. This ability is referred to as chord change negotiation, stressing its inherent problem-solving dimension.

  Note that continuity does not necessarily imply brownness, in the sense of [29], i.e. the sole use of small intervals. It rather implies that notes are glued together smoothly, and not made up of isolated elements or patterns, concatenated without care. For instance, the phrase in Figure 3 contains several large intervals but is perfectly continuous. It is straightforward to measure objectively and makes a relevant indicator of one's skills evolution.
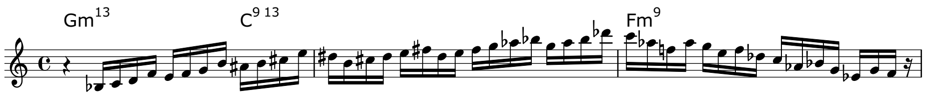
Figure 3. A virtuoso passage (152 bpm) in a chorus by John McLaughlin on Frevo Rasgado (1977), that contains several large intervals but is perfectly continuous.

Of course, in fine the main issue with automatic feedback does not consist in feeding the user with raw feature values but to compare their distribution with large corpuses of real solos played form professional musicians, such as the Weimar's Jazzomat solo corpus [1] or collections of Django Reinhardt's solo transcriptions.

This chapter focuses on the design on the full process-line implemented to provide an automatic feedback on the *in-scale rate* defined herebefore.

2. *Social Virtual Band*

Social Virtual Band (SVB) is an environment designed to provide the jazz student with a software support to record himself over realistic accompaniment and to store and manage his collection of solos over the Cloud.

With the Cloud architecture, the musician can retrieve and listen to his solos. He can possibly request automatic analysis and feedback from peers or from the algorithms. In other words, *Social Virtual Band reifies* the solos and builds a social network around this atomic piece of interaction, just like Facebook and SoundCloud work with text, images and audio tracks.

2.1. Architecture

The Lead Sheet DataBase (LSDB), introduced in Section 1.2, is the central element of *Social Virtual Band*, since the whole user experience gravitates around this collection of songs. Therefore, most of the intelligence involving the songs and the solos is provided by a back-end server, through web services. Figure 4 sums up the architecture of *Social Virtual Band*.

The jazz student interfaces with the back-end server through the client application. The application provides a simple interface for selecting a song among the LSDB collection. It then connects to the server to retrieve both a MIDI file with an accompaniment generated in a particular style, and the score of the chord progression. The application can then play the accompaniment while showing the score and recording the user's solo. The application holds the recorded solos, along with the generated accompaniment and some metadata, and sends the whole package to the back-end server, where it is stored.

The user can then connect to the web platform hosted by the remote server, and manage his collection of solos or follow the evolution of his playing skills.

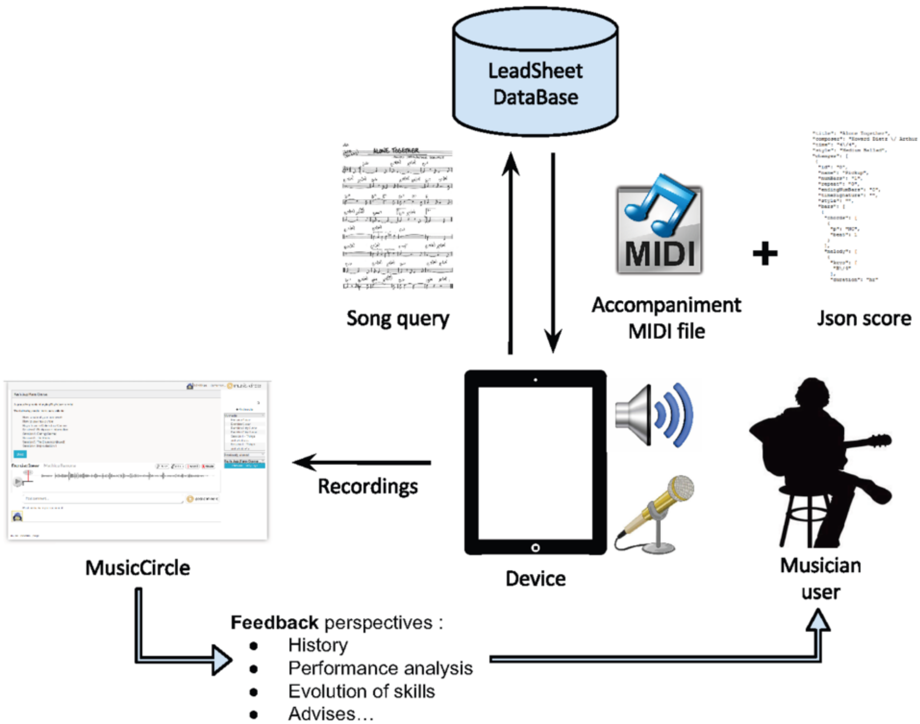

Figure 4. Architecture of the *Social Virtual Band* environment

2.2. Back-end server

The back-end server is based on an Apache Tomcat architecture[1] which provides various web services executed through Java code (designated as *servlets*).

The choice of using a web servlet is important in this context because it constraints us to define a clear, simple and yet versatile interface with the client, through the sole specifications of URL parameters. Indeed, each access to the server is performed with an HTTP request, that can be tested with any web browser, by typing the URL.

The web services involved in *Social Virtual Band*are implemented through so-called *Json Servlets*, i.e. servlets relying on the JSON (*JavaScript Object Notation*) for interfacing with the client. Figure 5 shows the list of servlets implemented in the *Social Virtual Band* server to communicate with the client application:

Database Listing Returns the list of songs that fulfill the input query, which can specify the composer, the song book source or the style.

Score Extraction Returns the chord progression of a lead sheet in the database, along with metadata associated with the song (time & key signature, style, tempo range, etc.). The client application uses that information to display the interactive score of the song.

[1] http://tomcat.apache.org

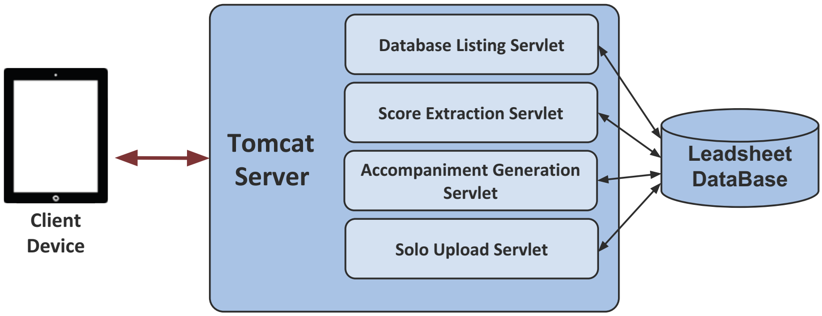

Figure 5. Architecture of the back-end Tomcat server

Accompaniment Generation This servlet is by far the most polyvalent provided by the server. It can generate a wide range of accompaniments by combining several algorithms developed by Sony CSL, detailed hereafter.

Solo Upload Retrieves the solo recorded with the application (with associated accompaniment and metadata), stores it into the database and performs the automatic processing to extract feedback.

2.3. Accompaniment Generation

The accompaniment is composed of three parts, following the usual convention for rhythmic sections in a jazz ensemble: drum, bass, and piano playing the voicing on the chord progression. In order to cope with the large variety of possible lead sheets to improvise on, it is natural to look for automatic accompaniment generators.

Several commercial systems propose automatic accompaniment systems (the most well-known being Band-in-a-Box and iReal Pro). Our architecture can cope with arbitrary accompaniment generators, taking as input a lead sheet and producing as output an accompaniment. Techniques vary greatly depending on the nature of the output. MIDI generators produce MIDI files which are then rendered by a embedded synthesizer, whereas audio generators generate directly audio files. We have experimented with both approaches.

Both the bass and piano part generation[2] are based on Markov constraints [20]. The Markov constraints is a technology for defining a Markov Model through a Constraint Satisfaction Problem (CSP). The CSP paradigm offers the possibility to specify hard constraints on sequences generated according to the Markov Model. This is particularly fitted for generating musical sequence, where many constraints apply for guaranteeing syntaxic and semantic consistency.

A MIDI generator was built using the technologies of Markov constraints, notably the meter constraint [23]. In a first step we have recorded accompaniments played by British jazz pianist Ray d'Inverno, over 20 jazz standards. In the generation step we take a given lead sheet as target and build a sequence of MIDI chunks (in this case, piano) that

[2]The generation of the drum part is still an on-going perspective, since the logic of chunks concatenation is very different when working on drums.

fit with the metrical structure of the lead sheet (the metrical location of chord changes). Then each chunk is adapted harmonically to fit with the target chords by changing the relevant pitches.

The accompaniment generated by the servlet is encoded into a MIDI file and sent to the client. The audio synthesis is performed in real-time by the client. Transmitting a symbolic MIDI file instead of raw audio dramatically reduces the volume of data downloaded by the remote client (typically, a 100 kB MIDI file can be equivalent to a 4 MB MP3 file encoded at 192 kbps bitrate).

2.4. Client Application

The client application provides a graphical interface for viewing and following a lead sheet and recording solos on generated accompaniments. It runs on Windows, MacOSX, Linux and iOS. The use of a portable device (smartphone or tablet) provides a very intuitive experience, since it relies on the embedded audio input and output devices, and requires no peripheral device.

2.4.1. User Interface

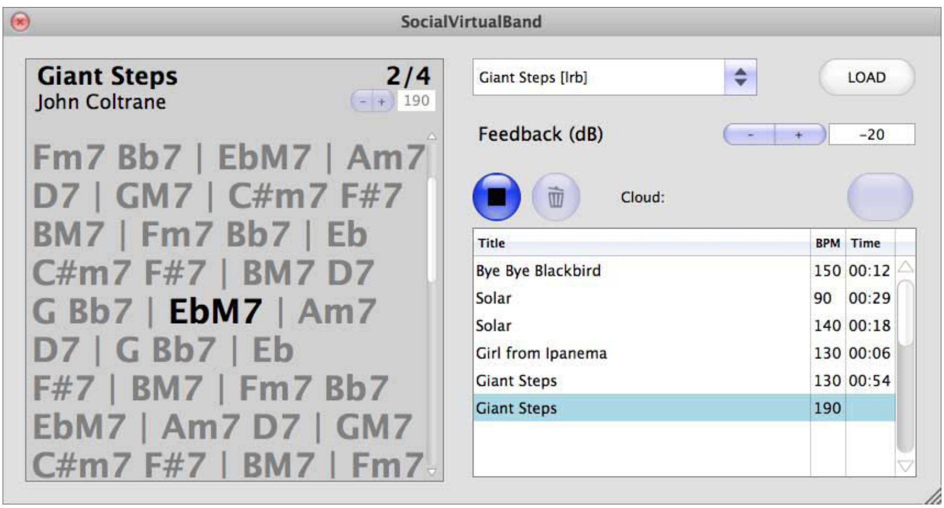

Figure 6. Graphical User Interface of the *Social Virtual Band* Client

Figure 6 shows a capture of the client user interface. The left panel shows the current lead sheet and follows the score during the recording process, while the right panel provides all the controls. The user can either select a pre-loaded song (from the upper choice list) or remotely long any song from LSDB by clicking the *LOAD* button. The lower table shows the list of solos recorded so far, and allows to replay them, and eventually upload them to the server.

2.4.2. Recording Process

The main issue dealt with is the recording process itself. Indeed, the application needs to play the accompaniment and to record the solo part (without the accompaniment) at the same time. And the recorded solo needs to be precisely synchronized with the accompaniment.

Figure 7 sums up the three possible recording configurations. The most straightforward configuration (a) consists in recording the solo while listening to the accompaniment played with the device loudspeaker. But this implies that the recorded solo track also captures the played accompaniment. This issue is fixed (b) by using headphones for the accompaniment. This way, the solo track is clean, but the configuration is less comfortable for the player, and the accompaniment might cover the hearing of his own playing. This latest issue is solved (c) by adding a feedback of the recorded solo to the played accompaniment. The *feedback* selector, visible in the GUI (Figure 6) lets the player modulate the feedback level to choose the right balance with the accompaniment. Of course this configuration only makes sense if the audio latency of the device is very low.

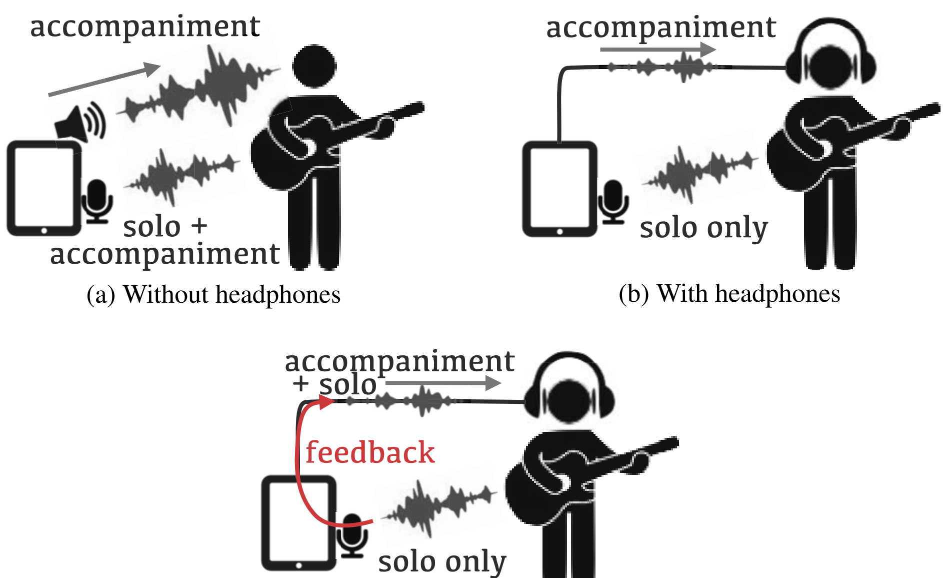

(a) Without headphones (b) With headphones

(c) Without headphones & feedback

Figure 7. Comparison of the different recording configurations when using the SVB client application

In the three situations depicted in Figure 7, maintaining the synchronicity between the generated accompaniment and the recorded solo is crucial for guaranteeing the alignment of the solo record with the chord sequence, and provide a precise temporal and harmonic analysis of the performance.

As illustrated by Figure 8, there is an inevitable latency (due to the sound device buffers[3]) between the instants the application generates a sound and the sound that is

[3]rather than the travel of sound in the air, marginal here considering the short distance between the player and the device.

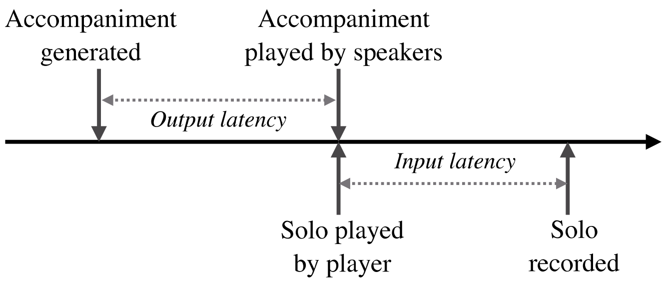

Figure 8. Output and Input latencies in the recording process

actually broadcast by the speakers (or headphones), and between the instants the player plays a note and that note is actually recorded by the application. These are respectively labeled as *output* and *input latency*, and are provided by the sound device itself. Supposing that the player plays simultaneously with the accompaniment, the recorded audio track has a delay that equals the sum of both latencies, and must be compensated.

In order to store the solo performance on the server without the accompaniment, the headphones solution, which is discussed in the preceding paragraph, is both simple and straightforward. Nevertheless, especially with respect to the playing comfort, it might be hindering and a set of headphones may not always be within reach. For this, Sony CSL in Paris has developed an alternative solution based on signal processing, using the frequency scale of the auditory system [10].

As a general rule, any cancellation algorithm works best when the solo in the recording is heard louder than the accompaniment. This can be easily achieved by means of a pickup attached directly to the instrument.

2.4.3. Upload to Server

As explained in previous Section, audio latencies are compensated to ensure a proper synchronization between the recorded solo track and the generated accompaniment. At the end of the recording process, the application holds both contents mutually aligned, along with the exact positions of chord changes, as depicted on Figure 9.

This information is sent to the remote server through a dedicated servlet that will analyze the audio track and use the aligned metadata to automatically extract relevant information.

2.5. Solo Analysis

2.5.1. Archiving

Each uploaded solo, when received on the server, is stored with the associated metadata. The server provides a Social Network (Figure 10) where each user can browse his own collection of solos.

The server keeps track of the whole history of solos previously recorded by the musician. It provides him an with statistical estimation of the evolution of his skills on a given track, for instance the evolution of the played tempo, as shown on Figure 11.

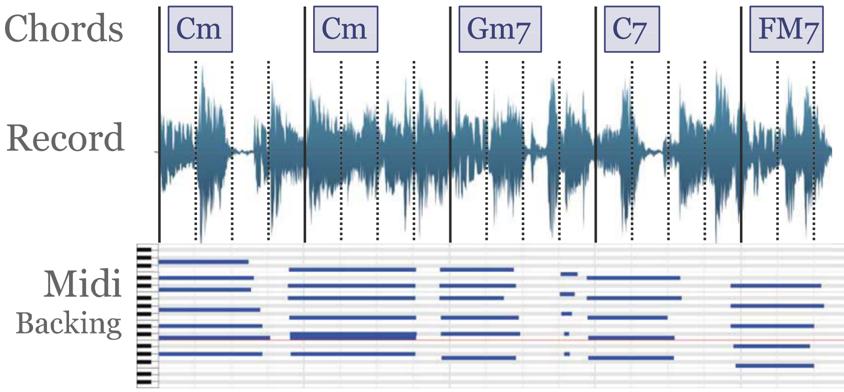

Figure 9. The recorded solo is sent to the remote server with properly aligned accompaniment and chord sequence.

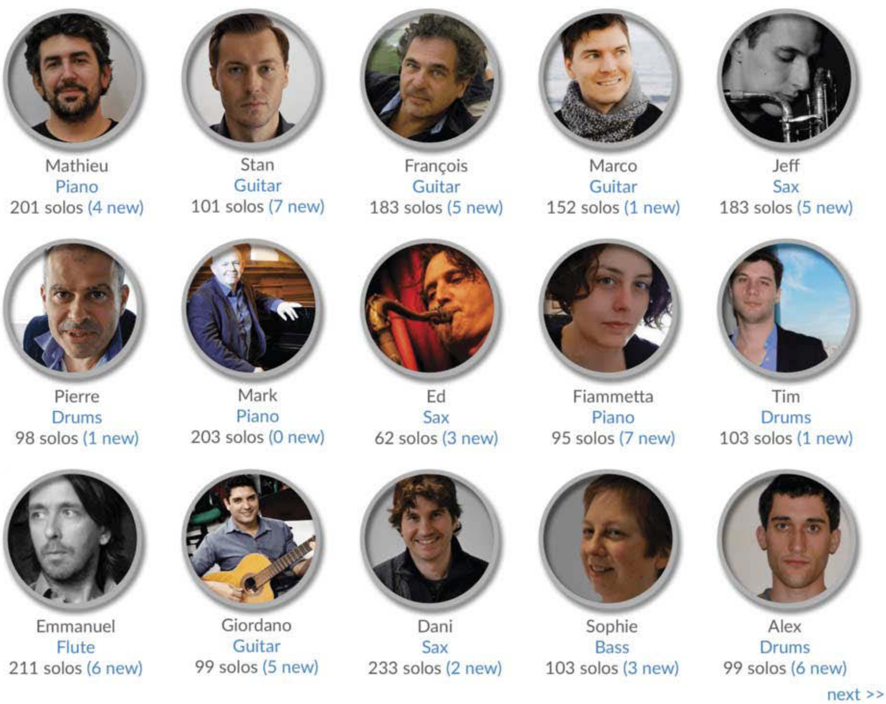

Figure 10. Interface of the Social Network coupled with the *Social Virtual Band* application.

A more detailed analysis is provided on each uploaded solo, based on an automatic transcription process.

2.5.2. *Transcription*

Many contributions in the literature cover the subject of audio transcription. The general problem of polyphonic transcription involving several musical instruments is complex [5] [30], often tackled with source-separation related methods, such as Non-Negative

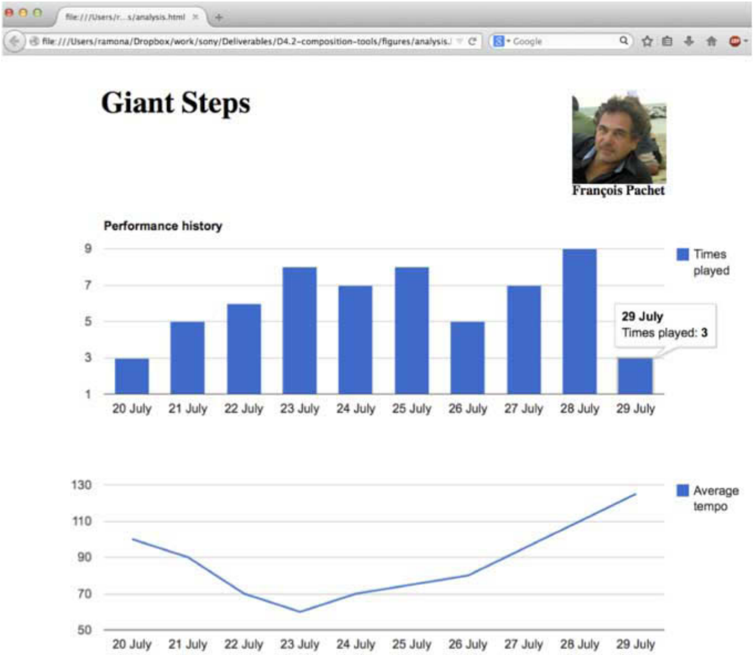

Figure 11. The *Social Virtual Band* web server provides statistical charts on the history of recorded solos. This figure shows the fluctuation of tempo used by the musician on the song Giant Steps.

Matrix Factorisation (NMF), which jointly estimates the dictionary and the decomposition of the audio signal into this dictionary [25] [2].

The polyphonic poly-instrumental problem is ill-posed, in the absence of (and even with) prior knowledge of the instrument timbre, because different instruments can play overlapping notes, and these notes usually share most of their harmonics (typically related by an octave or a fifth interval) because of usual musical assonance. A common way to avoid these ambiguities is to extract only the prominent melody [24].

A better-posed problem is the monophonic pitch estimation. The absence of overlapping between concurrent pitch harmonics simplifies the problem and results in a dramatic increase of state-of-the-art performance [6].

The *Social Virtual Band* use-case remains fairly simple because it is limited to mono-instrumental transcription of mostly-monophonic solos. The guitar and the piano allow the user to play chords in the solos, but we take the assumption that these polyphonic strokes are occasional and strictly vertical (as opposed to the horizontal polyphony of counterpoint). The system is currently using the guitar transcription algorithm developed by IIIA [16], but any other transcription algorithm could easily be plugged into the process line.

The melody extracted by the transcription algorithm is a sequence of notes, possibly polyphonic, bounded by start and end time values. In order to provide a symbolic representation (the score) of the transcribed solo, these temporal boundaries are quantized with a quantification step of 1/72 beat, which allows complex rhythmic divisions including sixteenth notes and triplets.

2.5.3. Harmonic Analysis

In order to evaluate the solo performances, we first analyze the lead sheets to extract the scales to be used for each chord. Obviously, real improvisation should not be perfect with regards to the underlying harmony, and "there is arguably some excess in the way improvisation is taught in jazz schools, focusing too much on the 'right' and 'wrong' notes" [Gilad Atzmon, personal conversation]. However, in our case we can use this information to produce an estimation of how far or close the improvisation is to the target harmony.

A lot of approaches have been used to analyze harmony (see, e.g., [18] for some references). In our case we use a simple dynamic programming approach, consisting in finding the harmonic analysis, for each chord label, which minimizes the number of modulations, i.e. scale changes. This process is performed through the following steps:

1. **Computation of possible harmonic analysis**
   For each chord label, we first compute (with the *MusES* library) the list of all possible harmonic analyses. A harmonic analysis is basically a scale (out of 3 possible scales types and 12 possible roots) and a degree. This process is described in [17]. We consider, for 3 basic scales (major, harmonic minor (hMinor) and melodic minor (mMinor), see Figure 12) all the scale-tone chords built by stacking up a number n of thirds, e.g.:

   Eb M7 ($n = 3$): [**I** of Eb Major, **VI** of G hMinor, **V** of Ab Major, **V** of Ab hMinor, **V** of Ab hMinor, **IV** of Bb Major, **IV** of Bb mMinor]

   C m7b5 ($n = 4$): [**VII** of Db Major, **VII** of Db mMinor, **VI** of Eb mMinor, **IV** of G hMinor, **II** of Bb hMinor]

2. **Definition of a dynamic programming problem**
   We introduce a cost function which assigns the following cost to a given transition between two harmonic analysis:

```
transitionCost(HarmonicAnalysis x, HarmonicAnalysis y) {
    if (x.getScale().equals(y.getScale()))
        return x.getDegree();
    else
        return 20 + x.getDegree();
}
```

3. **Computation of an optimal solution**
   In order to model the fact that tunes usually loop over themselves, we add the first chord at the end of the sequence.

As an example, Table 1 shows the possible analysis for each chord of the Giant Steps chord sequence. The dynamic approach described here produces the analysis shown in Table 2, which can be considered perfect on that example.

2.5.4. In-Scale playing rate

The harmonic analysis is not just an exercice for musicologist. It is in fact capital for the jazz musician because it defines the sequence of scales over which he can play without having to modulate. Indeed, the harmonic analysis defines a chord as the degree of a

B	**V** Eb hMin	**V** E Maj	**V** E hMin	**V** E mMin	**III** Gb Maj	**III** Gb mMin	**I** B Maj
D 7	**I** D Maj	**V** Gb hMin	**V** G Maj	**V** G hMin	**V** G mMin	**IV** A Maj	**IV** A mMin
G	**V** C Maj	**V** C hMin	**V** C mMin	**IV** D Maj	**IV** D mMin	**I** G Maj	**VI** B hMin
Bb 7	**VI** D hMin	**V** Eb Maj	**V** Eb hMin	**V** Eb mMin	**IV** F Maj	**IV** F mMin	**I** Bb Maj
Eb	**I** Eb Maj	**VI** G hMin	**V** Ab Maj	**V** Ab hMin	**V** Ab mMin	**IV** Bb Maj	**IV** Bb mMin
A m7	**VI** C Maj	**V** Db hMin	**IV** E hMin	**III** F Maj	**II** G Maj	**II** G mMin	**I** A hMin
D 7	**I** D Maj	**V** Gb hMin	**V** G Maj	**V** G hMin	**V** G mMin	**IV** A Maj	**IV** A mMin
G	**V** C Maj	**V** C hMin	**V** C mMin	**IV** D Maj	**IV** D mMin	**I** G Maj	**VI** B hMin
Bb 7	**VI** D hMin	**V** Eb Maj	**V** Eb hMin	**V** Eb mMin	**IV** F Maj	**IV** F mMin	**I** Bb Maj
Eb	**I** Eb Maj	**VI** G hMin	**V** Ab Maj	**V** Ab hMin	**V** Ab mMin	**IV** Bb Maj	**IV** Bb mMin
F# 7	**III** Db Maj	**III** Db mMin	**VII** Gb Maj	**V** Bb hMin	**V** B Maj	**V** B hMin	**V** B mMin
B	**V** Eb hMin	**V** E Maj	**V** E hMin	**V** E mMin	**III** Gb Maj	**III** Gb mMin	**I** B Maj
F m7	**IV** C hMin	**III** Db Maj	**II** Eb Maj	**II** Eb mMin	**I** F hMin	**I** F mMin	**VI** Ab Maj
Bb 7	**VI** D hMin	**V** Eb Maj	**V** Eb hMin	**V** Eb mMin	**IV** F Maj	**IV** F mMin	**I** Bb Maj
Eb	**I** Eb Maj	**VI** G hMin	**V** Ab Maj	**V** Ab hMin	**V** Ab mMin	**IV** Bb Maj	**IV** Bb mMin
A m7	**VI** C Maj	**V** Db hMin	**IV** E hMin	**III** F Maj	**II** G Maj	**II** G mMin	**I** A hMin
D 7	**I** D Maj	**V** Gb hMin	**V** G Maj	**V** G hMin	**V** G mMin	**IV** A Maj	**IV** A mMin
G	**V** C Maj	**V** C hMin	**V** C mMin	**IV** D Maj	**IV** D mMin	**I** G Maj	**VI** B hMin
C# m7	**VII** Db hMin	**VII** Db mMin	**VI** E Maj	**V** F hMin	**III** Ab hMin	**III** A Maj	**II** B Maj
F# 7	**III** Db Maj	**III** Db mMin	**VII** Gb Maj	**V** Bb hMin	**V** B Maj	**V** B hMin	**V** B mMin
B	**V** Eb hMin	**V** E Maj	**V** E hMin	**V** E mMin	**III** Gb Maj	**III** Gb mMin	**I** B Maj
F m7	**IV** C hMin	**III** Db Maj	**II** Eb Maj	**II** Eb mMin	**I** F hMin	**I** F mMin	**VI** Ab Maj
Bb 7	**VI** D hMin	**V** Eb Maj	**V** Eb hMin	**V** Eb mMin	**IV** F Maj	**IV** F mMin	**I** Bb Maj
Eb	**I** Eb Maj	**VI** G hMin	**V** Ab Maj	**V** Ab hMin	**V** Ab mMin	**IV** Bb Maj	**IV** Bb mMin
C# m7	**VII** Db hMin	**VII** Db mMin	**VI** E Maj	**V** F hMin	**III** Ab hMin	**III** A Maj	**II** B Maj
F# 7	**III** Db Maj	**III** Db mMin	**VII** Gb Maj	**V** Bb hMin	**V** B Maj	**V** B hMin	**V** B mMin

Table 1. Giant Steps chord sequence with possible analysis for each chord

B	**I**	of B Major
D 7	**I**	of D Major
G	**IV**	
Bb 7	**I**	of Bb Major
Eb	**IV**	
A m7	**II**	of G Major
D 7	**V**	
G	**I**	
Bb 7	**I**	of Bb Major
Eb	**IV**	
F# 7	**V**	of B Major
B	**I**	
F m7	**II**	of Eb Major
Bb 7	**V**	
Eb	**I**	
A m7	**II**	of G Major
D 7	**V**	
G	**I**	
C# m7	**II**	of B Major
F# 7	**V**	
B	**I**	
F m7	**II**	of Eb Major
Bb 7	**V**	
Eb	**I**	
C# m7	**II**	of B Major
F# 7	**V**	

Table 2. Result of the harmonic analysis for Giant Steps

tonic chord when the content of their scales are almost identical. The player can thus improvise on a D Major and still *sound* like a G as degree IV of D. The three scales used to characterize the modulations are defined in Figure 12.

For this purpose, the system uses the result of the harmonic analysis (i.e. the sequence of modulation scales) to verify that the transcribed solo fits with it. Each modulation covers a set of bars over which we compute the rate of played notes that belong to the scale. This so-called *In-scale rate* is provided as feedback to the user, along with the transcription of his solo.

Figure 13 shows an example of solo transcription, that also shows the chord sequence of the lead sheet, the harmonic analysis (each modulation is indicated by a trans-
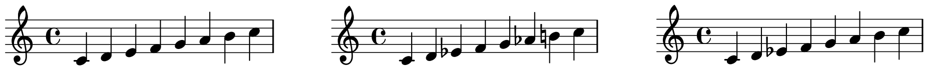
(a) Major (b) Harmonic Minor (c) Melodic Minor

Figure 12. Three basic scales of the harmonic analysis

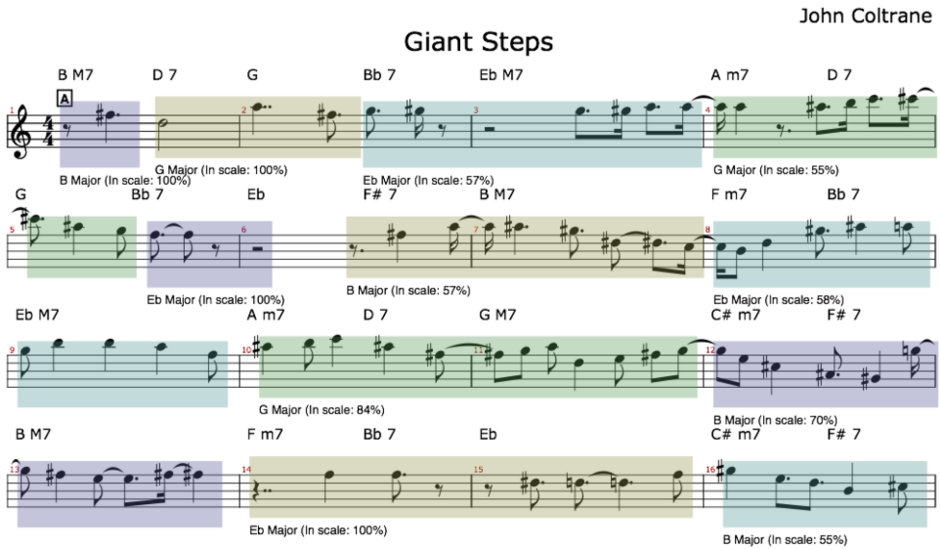

Figure 13. Result of the solo transcription process, displayed as a lead sheet score.

parent colored bar) and the *In-scale* playing rate, indicated for each modulation area.

3. Discussion

Social Virtual Band raises many issues related to jazz pedagogy and its social aspects. The system described here is a first step to provide the learning musician with an automatic feedback on his performance. Our focus here is to provide a measure of the student's skills estimated on a solo he performed. The in-scale rate, described herebefore, is an important measure that will soon be completed by the on-beat rate and the continuity measure, as defined in Section 1.3.

However, it is crucial to determine how this feedback does impact the evolution of the music student's skills. Such information will help us define future directions for music pedagogy support, based on automatic feedback. Future experiments are planned that will involve jazz students (from jazz schools) practicing improvisation with *Social Virtual Band* during a long period of a few months. Interestingly, the guitar practice video game Rocksmith 2014 claims that anyone can learn guitar with their system within only two months by playing one hour a day. Contrary to Rocksmith, *Social Virtual Band* is not designed as a substitute to the music teacher, but rather as a complementary support, that would be ideally nested within a social network.

Indeed, providing automatic feedback to the music student is a true innovation in the field of music pedagogy, since most existing systems tend to only provide an environment for practicing music. Nevertheless, no matter how relevant a machine can get, feedback from one's peers with always prevail, especially from trustworthy contacts such as a teacher or an experienced friend. The platform presented here is ideally suited for the emergence of such a social network, based on music practice, because it provides all the tools needed: a portable playalong system, a cloud server for archiving one's solo collection, and automatic tools to provide feedback on such collections. The social dimension will turn the solo into a social object, just like text, photos and audio tracks are today, thanks to existing social networks.

4. Conclusion

This chapter introduced the first system for providing support and feedback for practicing jazz improvisation. While most existing play-along softwares only provide prerecorded accompaniment to train on, *Social Virtual Band* extends that experience by collecting the history of recorded solos and embedding them inside a dedicated social network. As shown in this chapter, the collection and analysis of solos implies solving several technical issues, e.g. the precise synchronization of the recorded solo with the accompaniement, or the accompaniement cancellation in the recording process. The social ecosystem built around the solos, allows the musician to receive feedback on his skills, both from his community and from ad-hoc automatic analysis. We presented here an example of relevant automatic measure of the solos quality, based on the comparison of the transcribed notes with the expected scales deduced from harmonic analysis.

Nevertheless, two key issues remain open, that will be considered through long-term experiments involving a first community of users: what makes a good solo? and what kind of advice can make one improve his skills? Both questions will find answers from the analysis of social exchanges in tutor-teacher pairs.

5. Acknowledgements

This research is conducted within the Flow Machines project which received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 291156, as well as the Praise project (EU FP7 number 388770), funded by the European Commission under program FP7-ICT-2011-8.

We also thank and Daniel Martín and Timotée Neullas for their contribution to the web services.

References

[1] Jakob Abesser, Klaus Frieler, Martin Pfleiderer, and Wolf-Georg Zaddach. Introducing the jazzomat project - jazz solo analysis using music information retrieval methods. In *10th International Symposium on Computer Music Multidisciplinary Research, CMMR 2013*, pages 653–661, Marseille, Octobre 15-18th 2013.

[2] Nancy Bertin, Roland Badeau, and Emmanuel Vincent. Fast bayesian nmf algorithms enforcing harmonicity and temporal continuity in polyphonic music transcription. In *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, pages 29–32, New Paltz, New York, USA, October 2009.

[3] Michael A. Casey, Remco Veltkamp, Masataka Goto, Marc Leman, Christophe Rhodes, and Malcolm Slaney. Content-based music information retrieval: Current directions and future challenges. *Proceedings of the IEEE*, 96(4):668–696, April 2008.

[4] Jerry Coker. *Improvising Jazz.* "Simon & Schuster", 1964.

[5] Alain de Cheveigné. Multiple f0 estimation. In DeLiang Wang and Guy J. Brown, editors, *Computational Auditory Scene Analysis: Principles, Algorithms and Applications*. John Wiley and sons, Inc., September 2006.

[6] Alain de Cheveigné and Hideki Kawahara. Yin, a fundamental frequency estimator for speech and music. *Acoustical Society of America Journal*, 111(4):1917–1930, avril 2002.

[7] Robert R. Faulkner and Howard S. Becker. *Do you know ? The Jazz Repertoire in Action*. MIT PRess, 2009.

[8] Jon Gillick, Kevin Tang, and Robert M. Keller. Machine learning of jazz grammars. *Computer Music Journal*, 34(3):56–66, 2010.

[9] M. Gladwell. *Outliers: the Story of Success*. Back Bay Books, 2011.

[10] Stanislaw Gorlow, Mathieu Ramona, and François Pachet. SISO and SIMO accompaniment cancellation for live solo recordings based on short-time ERB-band Wiener filtering and spectral subtraction. 2015. submitted.

[11] Thomas Hedges, Pierre Roy, and François Pachet. Predicting the composer and style of jazz chord progressions. *Journal of New Music Research*, 43(3):276–290, 2014.

[12] M.B. Holbrook. *Playing the Changes on the Jazz Metaphor*. Foundations and trends in marketing. Now Publishers, Incorporated, 2008.

[13] Gary Kennedy and Barry Kernfeld. *The new Grove dictionary of jazz, vol. 1 (2nd ed.)*. New York: Grove's Dictionaries Inc., 2011.

[14] Mark Levine. *The Jazz Theory Book*. Sher Music Company, 1995.

[15] PG Music. Band-in-a-Box. http://www.pgmusic.com/.

[16] Tan Hakan Ozaslan, Enric Guaus, Eric Palacios, and Josep Lluis Arcos. Identifying attack articulations in classical guitar. In *Computer Music Modeling and Retrieval. Exploring Music Contents*, volume 6684, pages 219–241. Springer Verlag, 2011.

[17] François Pachet. An object-oriented representation of pitch-classes, intervals, scales and chords: The basic muses. In *Proceedings of Journées d'Informatique Musicale (JIM)*, Bordeaux (France), 1994. Université de Bordeaux.

[18] François Pachet. Computer analysis of jazz chord sequences: Is it Solar a blues ? In E. Miranda, editor, *Readings in Music and Artificial Intelligence*. Harwood Academic Publishers, 2000.

[19] François Pachet and Pierre Roy. Analytical features: a knowledge-based approach to audio feature generation. *EURASIP Journal on Audio, Speech, and Music Processing*, 2009(1), February 2009.

[20] François Pachet and Pierre Roy. Markov constraints: steerable generation of markov sequences. *Contraints*, 16:148–172, March 2011.

[21] François Pachet, Jeff Suzda, and Daniel Martin. A comprehensive online database of machine-readable leadsheets for jazz standards. In *Proc. ISMIR '13*, pages 275–280, Curitiba (Brazil), November 2013.

[22] Geoffroy Peeters. A large set of audio features for sound description (similarity and classification) in the CUIDADO project. Technical report, IRCAM.

[23] Pierre Roy and François Pachet. Enforcing meter in finite-length markov sequences. In *Proceedings of AAAI*, July 14-18 2013.

[24] Justin Salamon, Emilia Gomez, Daniel P.W. Ellis, and Gaël Richard. Melody extraction from polyphonic music signals: Approaches, applications and challenges. *IEEE Signal Processing Magazine*, 31(2):119–134, March 2014.

[25] Paris Smaragdis and Judith C. Brown. Non-negative matrix factorization for polyphonic music transcription. In *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, pages 177–180, New Paltz, New York, USA, October 19-22 2003.

[26] Mark J. Steedman. A generative grammar for jazz chord sequences. *Music Perception*, 2(1):52–77, Fall 1984.

[27] D. Sudnow. *The ways of the hand. The organization of improvised conduct*. MIT PRess, 1993.

[28] Technimo. iReal Pro. http://irealpro.com/.

[29] Richard F. Voss and John Clarke. "1/f noise" in music: Music from 1/f noise. *The Journal of the Acoustical Society of America*, 63:258–261, 1978.

[30] Chunghsin Yeh, Axel Röbel, and Xavier Rodet. Multiple fundamental frequency estimation and polyphony inference of polyphonic music signals. *IEEE Trans. on Audio, Speech and Language Processing*, 18(6):1116–1126, August 2010.

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-119

Chapter 9. Steps Towards Intelligent MOOCs

A case study for learning counterpoint

Katrien BEULS [a] Johan LOECKX [a]
[a] *VUB Artificial Intelligence Lab, Brussels, Belgium*

Abstract. Despite their overwhelming success, present-day Massive Open Online Courses are far removed from the student modelling capacities displayed by earlier Intelligent Tutoring Systems. Being mere content delivery tools, MOOCs typically lack a thorough assessment module as well as tools for personalising the learner's track. When learning music, particularly, these two properties are indispensable. This chapter surveys suggestions made by experts in the field of AI in education today towards the incorporation of ITS tools and techniques into MOOCs. Yet, more traditional student models and tutoring modules are not without shortcomings themselves and the real challenge lies in making active models of both the tutor and the student, which can be used to predict future learning tracks and set the right challenges. Agent-based tutoring systems offer an attractive framework for building such active tutor/student models. The proposed concepts are illustrated in the domain of music composition. A tutoring system has been implemented to teach students the craft of *counterpoint*, a commonly used strategy for learning polyphonic music composition. It is based on the theory of *flow* to keep students motivated and optimize learning.

Keywords. Adaptive Learning, Intelligent Tutoring Systems, MOOCs, counterpoint, music learning, student models, tutoring strategies, agent-based tutoring systems, music tutoring, online learning, theory of flow

1. Introduction

Since their appearance in 2011, Massive Open Online Courses (MOOCs) have become omnipresent in today's higher education landscape. Yet, although their rise is recent and their popularity large, the ideas that support these courses have been around for multiple decades (see [1] in this Volume). The first "teaching machine" was introduced in the fifties by the behaviorist B. F. Skinner in the form of an incremental mechanical system that would reward students for correction responses to questions [2]. The idea was later reinforced by the famous two-sigma problem that could show that student achievement in classroom interaction differs greatly from results obtained from individual tutoring [3]. If we transfer this idea into today's globally connected age, to what extent do participants in Massive Open Online Courses experience individual tutoring? Surely, in terms of evaluation, “in classes of 100,000 students, or more, instructors, no matter how many assistants they might have, are not going to be able to do the grading” [4].

MOOCs are mere content delivering tools today, replacing traditional university lectures, more than tools for assisting teachers in traditional classroom education. Despite their overwhelming success in terms of student numbers that these courses reach, the current first-generation MOOCs have two main shortcomings that are often mentioned by experts: (i) automatic assessment does not go beyond regular expression matching in simple self test questions at the end of each lecture segment and (ii) every learner follows the same learning path through the lectures, lacking any personalized tutoring. These two shortcomings are also reflected in the high number of dropouts (90-95%), which is often attributed to challenges similar to distance learning, such as time management. However, a comparative study showed that “MOOC students learned a bit more than students in a traditional university course, but less than students taught with an interactive engagement pedagogy” [5]. In sum, MOOCs as they are today are very useful in blended learning settings, where a human teacher incorporates MOOC material into their own lectures but currently less efficient in stand-alone education.

Two main paths are typically put forward to escape this deadlock situation in the online setting. First, by constructing knowledge in a collaborative way and by assessing each other’s work, students learn from each other. This kind of learning is sometimes referred to as collaborative or peer learning and found especially interesting in problem-solving kinds of domains, like design or music composition. Second, online courses form a testbed for adaptive learning techniques by means of intelligent automated tools. Indeed, Intelligent Tutoring Systems have focused for decades on building exactly such systems that perform automatic assessment based on extensive domain knowledge and student models. A student model can be defined as the set of beliefs that a tutor has about a student. These beliefs include the knowledge and skills of the student in the target domain, his learning preferences and other attributes. They can be inferred based on a student’s observable behavior: through his answers, actions or the results that he obtains.

Indeed, one would assume that such personalized education becomes crucial in big groups of learners. The large-scale data available in MOOCs hosts a huge potential for machine learning techniques to extract and generalize over learner patterns and offer individual learning tracks. This chapter tries to bridge the apparent gap between earlier intelligent tutoring systems (ITSs) and the (seemingly) abrupt rise of video-based MOOCs and argues for the need of an **active tutor and student model**, something which can only be achieved within an agent-based architecture.

This chapter is organised as follows. First, the current state of individual tutoring in MOOCS is reviewed. Next, a brief history on Intelligent Tutoring Systems is given to make the reader familiar with the basic terminology and architecture. The main contribution of this chapter, active tutor and student models, will be introduced in a fourth section and its capabilities demonstrated in the domain of music—more specifically counterpoint tutoring, a compositional technique. Finally, conclusions are drawn.

2. Individual tutoring in MOOCs

How well do MOOCs score in terms of facilities for individual tutoring? Many course designers have argued already that the real advantage of using MOOCs lies in their value in so-called *flipped classrooms* or hybrid education, where a regular lecturer relies on MOOCs only as content delivery for his course and uses it to spend more time on individual tutoring and discussions on the subject matter in physical interactions with the students in the classroom. A MOOC is then rather seen as one way of learning, which allows students to connect and collaborate by engaging in the learning process actively. Daphne Koller reported a higher-than-usual attendance in her Stanford courses that are taught this way: "We can focus precious classroom time on more interactive problem-solving activities that achieve deeper understanding—and foster creativity" [6]. Such studies point to the importance of teachers as individual mentors who can debug students' thinking and "honestly be enthusiastic when they excel" [7].

Apart from improving the quality of face-to-face time in lectures, the real question is whether MOOCs can be used as stand-alone tools in distance education in the way intelligent tutoring systems were thought to function? A systematic comparative study by Judy Kay and her colleagues at the University of Sydney across a sample of major massive open online course platforms revealed that "all of the systems currently have only rudimentary facilities to capture learner activity data for analysis" [8], meaning that offering opportunities for individual learning paths is not on the agenda of current MOOC designers. The student can see "rather simple information about their marks and progress" [id.]. The researchers therefore rightfully conclude by saying that "here is a place where there is exciting potential to introduce Artificial Intelligence & Education (AIED) tools and techniques into MOOCs" [8].

MOOCs certainly offer new opportunities for individual tutoring when they are used in blended learning settings where teachers can inspect every student's individual progress on certain exercises and intervene if needed by offering targeted feedback on the components the student is struggling with or teaming up stronger with weaker peers [9]. Still, automatized individual tutoring with richer evaluation models to measure student engagement are not yet fully realized. This is not the least the case in learning music, a discipline that involves many different skills that cannot be assessed with simple multiple choice questions and in which the personal nature of a learner's track is primordial. Yet, the effective use of online and intelligent technology has been underresearched in music as well [10]. Their absence has contributed to the general feeling of disappointment towards MOOCs [11]. To understand how this aspect could be improved, the following section situates MOOCs within the larger history of Intelligent Tutoring Systems, in which individual tutoring through the use of student models plays a prominent role.

3. A brief history of Intelligent Tutoring Systems

Although today a well-established concept, Intelligent Tutoring Systems (ITSs) have gone a long way since the first breakthroughs in the early seventies that incorporated AI techniques into programmed instructions. These early advances allowed for (i) alternative representations of content, (ii) alternative paths through material and (iii) alternative means of interaction. Much of the research into expert systems turned out to be useful for representing expert (tutor) knowledge and building student models. In the eighties, and still very much so today, the main research questions of the field of ITS could be formulated as follows [12]:

- What is the nature of knowledge, and how is it represented?

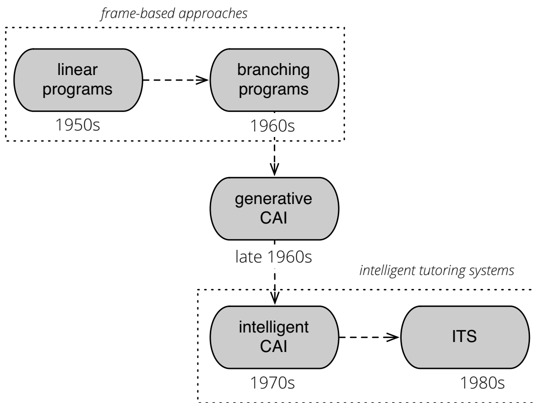

Figure 1. The history of Intelligent Tutoring Systems

.

- How can an individual student be helped to learn?
- Which styles of teaching interaction are effective, and when should they be used?
- What misconceptions do learners have?

3.1. *From frame-based approaches to ITSs*

The first attempts to build tutoring systems were all frame-based, where most frames contained simple questions (fill the gap exercises, selecting the correct answer, etc.). Such tutoring systems proceeded to present the next frame regardless of the accuracy of the student's response. It was therefore nothing more than a programmed textbook, completely lacking any individualization. In the 1960s, Crowder tried to overcome this major shortcoming as he introduced the notion of branching programs. Although still having only a number of fixed frames, these programs no longer ignored student's responses but the system could comment on a student's response and use it to choose the next frame [13].

Generative Computer-Assisted Instruction (CAI) was launched in the late 1960s. The idea of generative CAI was that a computer could generate teaching material automatically. One of the main advantages was that memory usage could be considerably reduced by this technique, since the frames did not have to be saved as such. However, this approach remained restricted to drill-type exercises, in which the learner model consisted of nothing more but an integer. Uhr and his collaborators [14] implemented a series of systems which auto-generated problems in vocabulary recall and arithmetic, two domains that presumably require drill and practice types of exercises. The sophistication in their systems was situated in the task-selection mechanism, which ensured the exercise level to be adapted to the student's overall performance.

It was Jaime Carbonell's mission to put Artificial Intelligence into CAI, meaning that the computer should have a representation of what is being taught, to whom and how [15, 16]. He developed SCHOLAR, a tutoring system for teaching Latin-American geography. SCHOLAR helped students enhance their knowledge by (i) solving problems at a certain level or by (ii) involving them in discussions with the computer in a more interactive way.

Although there is no sharp boundary, in the 1980s intelligent CAI was replaced by Intelligent Tutoring Systems (ITS), which try to extend the domain of applicability, power and accuracy of CAI systems [17, 18, 19, 20]. Figure 1 summarises the different steps in the history of Computer-Aided Instruction until the arrival of ITS. Examples of early ITS include the Pittsburgh Urban Math Project (PUMP) algebra tutor [21] and the SHERLOCK control panel [22], used to train Air Force techniques to diagnose problems that might occur. Another early system is GUIDON [17, 23], which was the first intelligent tutor based on an expert system. GUIDON was also the first program to teach medical knowledge.

3.2. The general ITS architecture

Current Intelligent Tutoring Systems have a standard architecture with a number of components that are each responsible for a specific function. The components can best be explained according to the knowledge type they encode, which results in the following four types:

1. **Domain knowledge** (how experts perform in the domain): definitions, processes or skills needed to multiply numbers (e.g. the AnimalWatch tutor), generate algebra equations (e.g. PAT tutor), etc.;
2. **Student knowledge** (how to reason about student knowledge): stereotypic student knowledge of the domain and information about current student (time spent on problems, hints requested, correct answers, preferred learning style, etc.);
3. **Tutoring knowledge** (encoding reasoning about the feedback): either derived from empirical observations of teachers or enabled by technology (simulations, animated characters);
4. **Communication knowledge**: includes graphical user interfaces, animated agents, dialogue mechanisms.

The domain knowledge module (expert knowledge), the student model module (student knowledge) and the tutoring module (tutoring knowledge) are interconnected in the main architecture of an ITS. Communication knowledge is incorporated by a user interface module that mediates between the student input and the tutoring module (Figure 2). Because the communication knowledge is often included in the tutoring module, the remainder of this section discusses the three main interconnected modules in the ITS architecture: expert knowledge, student knowledge and tutoring knowledge.

3.2.1. Expert knowledge

Domain models interact very closely with the student model: they are the first step in representing the expert knowledge. They can generally be divided into three categories of complexity: (i) problem solving (mathematics problems, Newtonian mechanics), (ii) analytic and unverifiable domains (ethics, law) and (iii) design domains (architecture, music composition). There are two main axes in the classification of domain models: a first one ranging from simple to complex and a second one ranging from well-structured to ill-structured. Category 1 models represent expert knowledge in the field of arithmetic

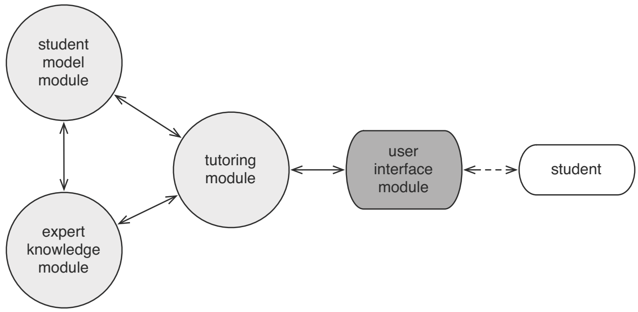

Figure 2. The basic architecture of an Intelligent Tutoring System consists of three main modules: expert knowledge, a student model and a tutoring module.

and other well-defined domains (well-structured, simple). Category 3 represents the other side of the axes: complex and ill-structured domains such as the knowledge needed to build an ITS for architecture tutoring. Finally, Category 2 contains qualitative representations of expert knowledge for fields such as language or music, which are halfway on both axes [24].

Learning music involves a whole range of learning activities: from music history over mastering an instrument to interpreting a piece and composing in a particular style, requiring different kinds of cognitive skills. While more academically oriented courses could be assessed and taught in a more traditional way, teaching students how to compose requires advanced tutoring and close guidance. In an online tutoring system, the models that represent this knowledge thus need to reflect this knowledge closely. To teach students the craft of classical composition, for example, a common approach is to teach the *practice of counterpoint*, which consists of a set of do's and don'ts in polyphonic music writing. It could be classified as Category 1 due to the formal and well-defined nature of rules that constrain harmony and melody. Still, just checking student's work for correctness is unsatisfactory given the complexity of the task and modelling the knowledge including all student misconceptions, is a difficult and time-consuming task that often needs to be carried out by hand.

3.2.2. Student knowledge

A student model can be defined as the set of beliefs that a tutor has about a student. These beliefs include the knowledge and skills of the student in the target domain, his learning preferences and other attributes. They can be inferred based on a student's observable behaviour: through his answers, actions or the results that he obtains. Traditional ITSs keep track of a student's performance based on a series of pre-set learning objectives, such as a range of grammatical phenomena in the target language or vocabulary items covering the learning situations that the student has selected.

To improve the student modelling enterprise, some tutoring systems allow their students to inspect and control the student model [25]. Student models with this property are called *open learner models* (OLMs). They can contain simple overviews of knowledge (such as a skill meter) or a more detailed representation of knowledge, concepts etc. Park Woolf [24] lists several motivations for the use of open learner models, such as (i) the student has the right of access to and control over his personal information;

(ii) the student can potentially correct the learner model; (iii) the frequent asymmetric relationship between the student and the tutor can be resolved; and (iv) OLMs stimulate reflective learning in the student.

Researchers in ITS tend to classify their student models according to three main dimensions. The first one covers the input that the system receives, while the remaining two are structural properties of the student profile. Van Lehn [26] refers to them as bandwidth, target knowledge type and the differences between student and expert:

1. *Bandwidth* refers to the amount and quality of the input that the diagnosis component receives about what the student is doing or saying. From this input, the tutor must infer what the student is thinking and believing [26].
2. *Target knowledge type*. Because a good student model can in practice solve the same problems as a real student would be able to solve, it can be used to actively predict the real student's answer. To solve these problems the model needs "*some kind of interpretation process that applies knowledge in the student model to the problem*" [27]. Depending on whether we are dealing with procedural or declarative knowledge, a different interpretation process is required.
3. *Student-expert differences*. The knowledge of a student is usually regarded as the background knowledge of a student modeling system. Student knowledge always needs to be understood in relation to an expert model that can provide explanations on the correct way(s) to solve a problem. To compare student and expert or tutor knowledge, most ITSs claim to use the same knowledge representation language for both [26]. However, reality is often different. Due to economy and other implementation issues, the student model is often a copy of the expert model plus a collection of differences: missing concepts (knowledge that the student does not yet have) and misconceptions (knowledge that the student has that the tutor does not).

3.2.3. Tutoring knowledge

A tutoring model has two main functions, which are mirrored in the basic tasks of instruction, namely *to stimulate and evaluate learning*. ITS resarch has mainly addressed these functions separately [28] and sometimes together as in the ASSISTment system [29, 30, 31], which combines 'assistance' and 'assessment'. A tutoring model thus needs to decide on *when* and *how* to intervene and it is responsible for content planning of what to teach next. The question of when and how to assist the learner is "the fundamental dilemma of tutoring" [32, 33, 34].

Assisting and tutoring the learner can further be divided into two sub-functions [34]: "cognitive diagnosis, defined as the detection of the sources of errors, and the selection of tutoring or remediation strategies". Recent developments in automatically learning the learner's affective states [35, 36, 37] have increased the complexity of reasoning about optimal tutoring decisions.

Tutor's decisions are often reflected in the different forms of interaction that the tutor has with the learner. Typical forms of interaction include socratic dialogs, hints, feedback from the system, etc. A human teacher typically uses six types of feedback [38, 39, 40]: explicit correction, recasts, clarification requests, metalinguistic feedback, elicitation, repetition or any combination of these. These interactions usually occur through the user interface module, that connects the student with the tutoring module (see Figure 2). The user interface often includes a dialogue system for interacting with the student. This type of conversational interaction is particularly useful when the learner's answer is incomplete. Because tutors usually have an approximate sense of what a student knows, it "appears to be sufficient to provide productive dialogue moves that lead to significant learning gains in the student" [41].

4. Introducing active tutor and student models

4.1. State-of-the-art

Intelligent Tutoring Systems today work with static student models that keep track of a student's performance on certain predefined knowledge domains by counting scores and comparing these to averages. If the technique of a student model is to be included to make MOOCs more intelligent, we should consider a more dynamic student model that can actually function as a real model of the actual student and predict future behavior.

In Artificial Intelligence, a promising way to introduce such dynamic models is to make use of agents that can take on the role of learners and tutors. Such agents are autonomous entities that pursue their own goals and learn according to the outcome of its own or other agents' actions. Indeed, multi-agent systems have sometimes been considered as good candidates for building basic Intelligent Tutoring Systems infrastructures as they fulfil all the necessary requirements [42]:

(i) they are made of different interconnected, complex components;
(ii) they provide multiple, different and complementary services;
(iii) each of their components is functionally autonomous; and
(iv) they are equipped with specific knowledge structure and reasoning mechanisms.

Agents are thus often decomposed by their function in the teaching and learning process, with for instance one evaluation agent, one modeling agent, one recording agent, one student agent, etc. [43]. Moreover, the usefulness of agent technology in intelligent education systems is their contribution to make these systems adaptive, able to learn and dynamic by providing dynamic adaptation of domain knowledge and of behaviour of individual learners ([44], cited by [43]). Pedagogical agent-based systems are often used to monitor a particular project and enhance communication between members of a group [45]. Some researchers have designed agents for every course unit [46] or assigned a new agent to a specific learning topic [47].

4.2. Active student & tutor agents

We suggest to abandon such distributed systems in favour of a more holistic agent-based tutoring system with only two agents: one that models a tutor and one that simulates a learner. The latter can then function as *an active student model* that can run and try out solutions in parallel to predict a student's behaviour. These predictions can then be used to set the right challenge level, select the next exercise and suggest corrections to the real student. We will reuse concepts and findings of a recent PhD dissertation on language tutoring with such an agent-based architecture [48]. Similarly, the two agents that form the backbone of an agent-based counterpoint tutoring system are the music agent (ex-

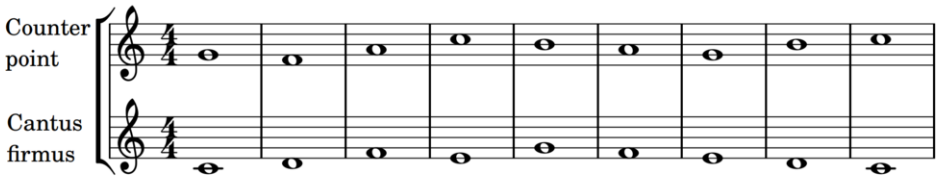

Figure 3. Example of a correct first species counterpoint piece. Given the cantus firmus (CF), a student is asked to compose the counterpoint voice, that should obey harmonic, melodic and motion constraints.

pert composer) and the student agent (music pupil). Both agents share the same architecture: a set of music rules, a processing engine and a meta-level architecture to capture inconsistencies in the student's composition of the counterpoint voice. Yet, although the components of the teacher and student agent are homologous, the realization of these components is not identical. The teacher agent represents an ideal composer in counterpoint whose musical skills also allow him to correct erroneous utterances of others. The student agent does not yet master all the counterpoint rules that are needed to be fully expressive in his compositions.

A music agent can be extended with a tutoring strategies component and a student profile component to become a fully-fledged tutor agent. These components personalize the tutoring process by keeping essential information about the student that is constantly being updated so that tutoring can be personalized to better fit the motivations of the individual student. This approach is particularly interesting in the domain of music, where learner's paths are very personal and the rules to be learnt exhibit complex interactions so that keeping track of which rules a student masters, is crucial.

5. Illustration in the domain of music

5.1. The study of counterpoint

The study of counterpoint attempts to express the properties of melodious polyphonic music, by investigating how the individual voices are formed and interact with each other. It has been a very important historical effort to capture the style of Renaissance music in which the independence of voices and harmonious polyphony is central. Still today, however, it is an indispensable tool to teach students the craft of classical music composition. A *counterpoint exercise* consists of a given *cantus firmus* (CF) or monophonic melody, on which a student has to compose a second melody to (called the *counterpoint voice*, CP). Of course, not all possibilities of notes are allowed (in this case a random melody would suffice), and here come the counterpoint rules into play. Figure 3 gives an example solution of such an exercise. The composed melody, should obey:

(a) harmonic or vertical (spanning two voices),
(b) melodic or horizontal (concerning one voice) and
(c) motion rules (relative movement of two voices).

As these rule are interacting in different ways, solving a counterpoint exercise can be seen as solving a complex sudoko. One discriminates between hard constraints (absolute "no-go's") and soft constraint (discouraged). Besides constraints, there are also guidelines

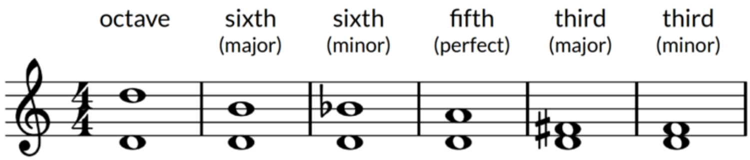

Figure 4. The study of counterpoint is centred around “rules” that limit the compositional freedom of a composer. They are intended to maximise the independence of voices while meanwhile preserving the harmonious nature. The most basic rule of counterpoint, pictured above, specifies which harmonic intervals are allowed and which are forbidden.

for creating ‘better sounding’ melodies, for example that a melody should be smooth and consist of the least number of skips available. The rules and guidelines collectively define the polyphonic musical style of the Renaissance era (1400-1600 AD). Because the ‘freedom’ of compositional expression has been limited on purpose in counterpoint, it helps pupils to cope with the many and complex interactions that occur constantly in polyphonic music.

Figure 4 depicts a basic harmonic rule in standard musical notation, that says that the possible *harmonic interval* (that is, distance between two notes with regards to pitch) are restricted: only octaves, minor/major sixths, perfect fifths and minor/major thirds are allowed. Of course there are many more and more complex rules that involve the two voices and movement of the melodic lines. The difficulty of counterpoint lies in the fact that the rules often conflict with each other so that solving one error leads to violation of another rule. It leads us too far to explain all this in detail and we refer readers interested to know more about the musical aspects of counterpoint to excellent literature on the topic [49, 50].

Also from a computational point of view, counterpoint is interesting as the “rules” to be learnt are quite formalised and the interactions not too complex. For this reason, it has been an interesting case study for computational representations in the past, using formalisms in the domain of feature spaces, fuzzy logic or constraint programming [51, 52, 53]. Though these approaches yield satisfying results, they lack musicality, explanatory power and mechanisms to engage them in tutoring activities as the musical knowledge can not be exploited.

In the following sections we will outline the basic operation of our proposed tutoring system for teaching counterpoint. We will look in more detail at the two-agent architecture, at the musical grammar to analyse a given composition and at how the tutor agent diagnoses errors and repairs them.

5.2. *Tutoring based on the theory of “flow”*

Figure 5 depicts the overall process of the proposed Intelligent Counterpoint Tutoring System, based on the concept of *flow*. Flow theory, introduced by Mihalyi Csikszentmihalyi [54], describes the experience of intrinsically motivated people when they are deeply engaged in an activity. During this time, people are in such a state of extreme concentration that they forget about time and the world around them. It is believed that learning is optimized in this situation and the primary goal of the process pictured in Figure 5 is thus to keep students in a state of flow. One of the groundbreaking aspects

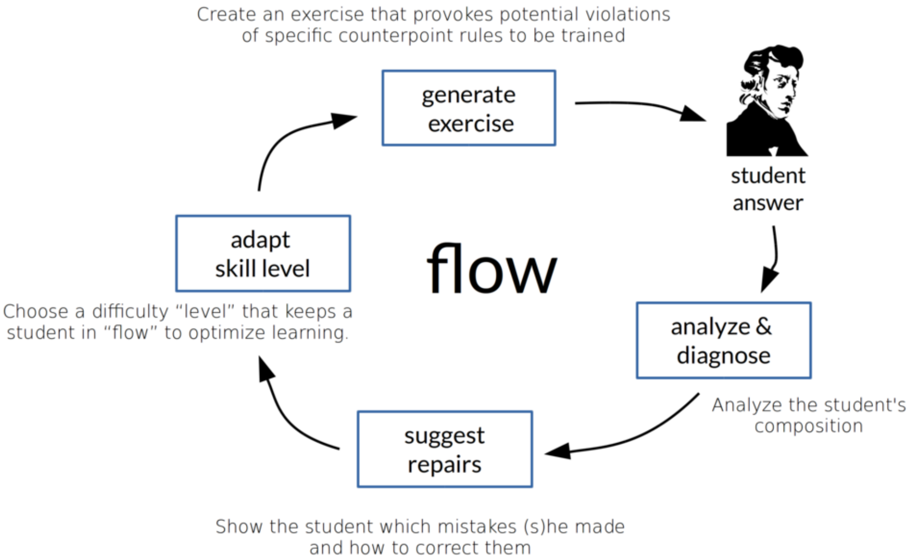

Figure 5. The challenge level is adapted to match the student's capabilities and skills to keep them in a state of flow and optimize learning. In order to achieve this, the tutoring agent should be well aware of what knowledge a student masters by keeping the student model up-to-date. The virtual tutor analyses the composition of a student and can suggest repairs when needed.

of this theory is the fact that the distinction between teacher and learner is diffused and students are put central to the learning process [55].

Flow occurs when a person perceives the challenges and the skills brought to it, as both balanced and (slightly) above average. For example, if an exercise is too difficult, a student may become anxious or frustrated which not only hampers learning but is also detrimental to motivation. If, on the other hand, the exercises that are presented are too easy, a student may become bored and lose interest. Finding this balance is one of the challenges of our tutor agent. As a consequence, it is essential that the tutor continuously gauges the student's skill level and keeps the student model in sync by tracking in detail which knowledge the student masters and which not.

In our specific case, the student model is an active one and corresponds to the (executable) set of counterpoint rules that the student masters. Because this knowledge can be instantiated, the tutor agent can *simulate* which response(s) a student will/might give. If there are any discrepancies observed, the student model is adapted to reflect accordingly to match the student's skill set. When a student provides an answer to an exercise, the expert tutor (who masters all counterpoint rules), analyses the piece and diagnoses which rules have been violated and suggests repairs. We will now look in more detail into each of these steps, starting with the musical analysis when a student submits the solution to an exercise.

5.3. Analysing the student's answer

To encode musical knowledge, we have opted for a grammar-based representation of counterpoint rules. Figure 6 shows an example initial structure before analysis. To build this grammar for musical composition we use the Fluid Construction Grammar frame-

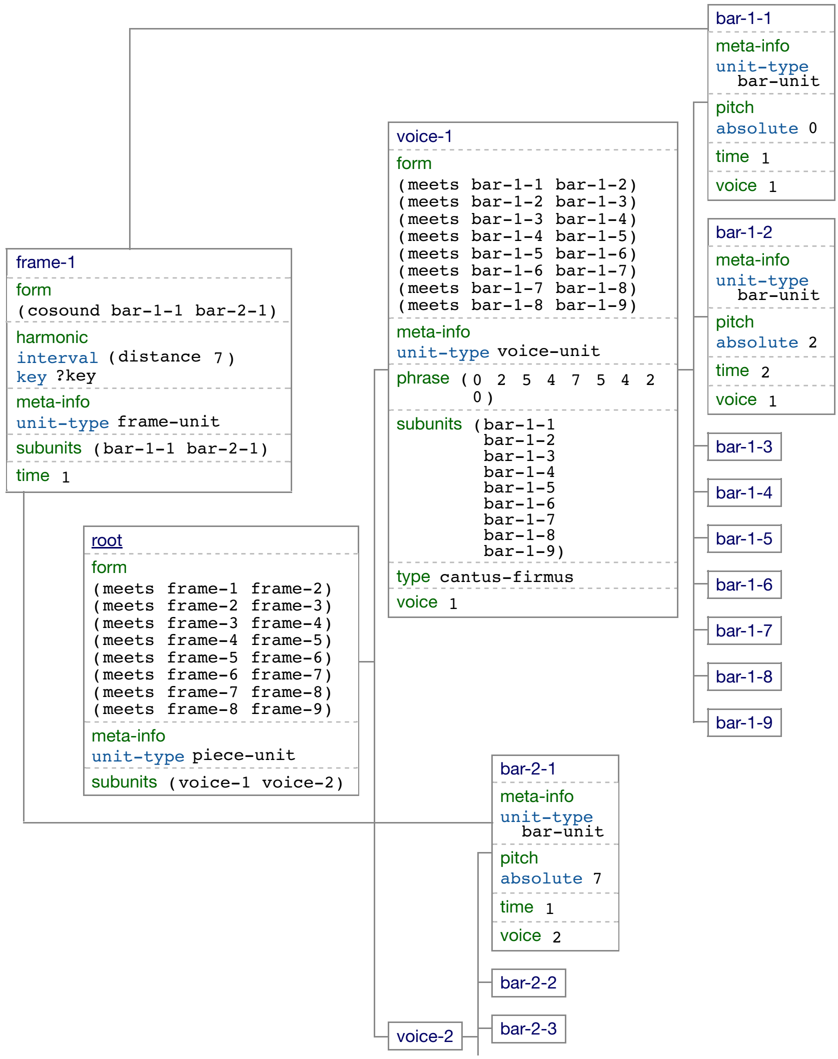

Figure 6. Partial initial transient structure for the example counterpoint exercise with selected voice, frame and bar units expanded. Bar units have two parent units: the voice to which they belong (horizontal view) as well as the frame (vertical view).

work, a computational formalism that is based on the theory of Construction Grammar within linguistics [56, 57]. In construction grammar, the main data structure is a construction, which is a mapping between meaning and form through semantic and syntactic categorisations. Constructions are implemented as feature structures with a conditional and a contributing pole. In grammatical processing, constructions contribute to a transient feature structure that is being built up from scratch, starting from a conceptualization in formulation and from an utterance in comprehension. Information that is in the conditional pole functions as a requirement that needs to be satisfied before a construction can contribute new features or hierarchy to the transient feature structure.

For example, a lexical construction for “goat” has the following conditional slots (on the right-hand side of the arrow): semantic predicates (conditions in formulation) and strings (conditions in comprehension). In formulation, this construction will contribute the “goat” form and features such as `referent`, `lex-cat`, `number` and `sem-cat`. In comprehension, the construction contributes the meaning predicates as well as the same contributing features on construction’s the left-hand side.

$$\left[\begin{array}{l}\textbf{?word-unit}\\ \hline\hline \text{referent: ?ref}\\ \text{lex-cat: noun}\\ \text{number: singular}\\ \text{sem-cat:}\\ \quad\{\text{animate, non-human, four-legged}\}\end{array}\right] \leftarrow \left[\begin{array}{l}\textbf{?word-unit}\\ \hline\hline \text{\# predicates: }\{\text{goat(?ref)}\}\\ \hline \text{\# form: }\{\text{string(?word-unit, “goat”)}\}\end{array}\right]$$

Similarly, musical rules can be represented as constructions. For the moment, we will ignore the controversial topic of ’meaning of meaning’ and consider all counterpoint rules as being purely syntactic. Now instead of features such as `form`, `referent` or `sem-cat`, a construction contains information about the `pitch`, `interval`, `motion`, etc. of a note or a musical phrase. Instead of an utterance, the grammar engine will now receive a musical piece with two voices, the first one being the cantus firmus and the remaining one the counterpoint voices. For simplicity sake, we now only consider *first species* counterpoint exercises, also called *note against note*, in which only one note is played at the same time as the cantus firmus, as shown in Figure 3. In this simple case, a melody can be represented by a straightforward list of notes , in which tone pitches are represented by a single number from 0 to 12, representing the relative pitch in a chromatic scale (an octave contains 12 semitones in Western music). The list-based representation of the piece discussed above, is included here:

```
((0 2 5 4 7 5 4 2 0)
 (7 5 9 12 11 9 7 11 12))
```

Figure 6 shows the initial transient structure when analysing the example piece. There are melodic units such as `voice-1` containing the full phrase of the cantus firmus and harmonic units such as `frame-1` for the cosounding bars 1-1 and 2-1 with an interval distance of 7 semi-tones (see pitch features in respective bar units). It is on this initial feature structure that individual constructions will work.

Constructions can then focus on any part of the transient structure: certain bar units within the same voice, a frame unit connecting two bars vertically, etc. or even units that have not been built yet such as a melodic motion unit, combining all adjacent notes within a single motion movement (see Figure 7). Similar to the lexical construction shown above, a harmonic construction such as the octave construction consists of a conditional pole with features that have to be present (here a frame unit with harmonic interval distance of 12) and a contributing pole (harmonic interval categories octave and consonant).

$$\left[\begin{array}{l}\textbf{?frame-unit}\\ \hline\hline \text{harmonic:}\\ \quad\text{interval:}\\ \quad\quad\text{category: }\{\text{octave, consonant}\}\end{array}\right] \leftarrow \left[\begin{array}{l}\textbf{?frame-unit}\\ \hline\hline \text{meta-info:}\\ \quad\text{unit-type: frame-unit}\\ \text{harmonic:}\\ \quad\text{interval:}\\ \quad\quad\text{distance: 12}\end{array}\right]$$

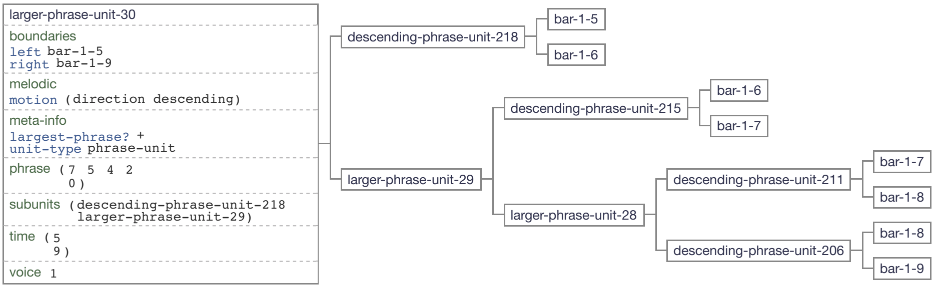

Figure 7. A descending motion phrase ranging from bar 1-5 to bar 1-9 of the example piece. Only the uppermost unit is fully expanded here. It groups a low-level descending phrase unit and a higher order phrase unit. The motion construction thus always looks at two units with the same melodic motion direction and unites them into a phrase unit, containing the full phrase, time information and the direction.

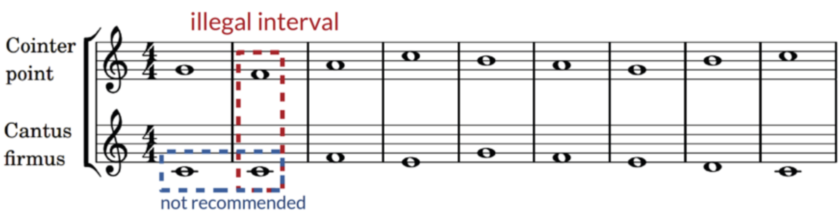

Figure 8. Example of a first species counterpoint piece with a violation of the "legal intervals" rule in the second bar.

The current grammar contains 29 constructions, divided into harmonic, melodic and harmonic-melodic ones. As it goes beyond the scope of this paper to describe the full grammar in detail, we'll simply make an inventory of the constructions of the tutor grammar that are used to analyse musical pieces.

5.4. Diagnosing the student's answer with regard to counterpoint rules

Once constructions have done their work to analyse the musical piece and build a transient structure that highlights relations between individual notes on a melodic and a harmonic level, the tutor uses a set of diagnostics that contain specific counterpoint rules to diagnose particular deviations. This section highlights some examples of such diagnostics for three types of counterpoint rules: harmonic, melodic and harmonic-melodic (motion). All diagnostics make use of the final transient structure that results from the application of the musical constructions. A diagnostic will either return one or more problems or nothing at all. The problems will then trigger repair strategies that launch a search process to find a satisfying solution for every problem. Let us consider the following erroneous counterpoint piece, with its musical notation shown in Figure 8.

```
((7 5 9 12 11 9 7 11 12)
(0 0 5 4 7 5 4 2 0))
```

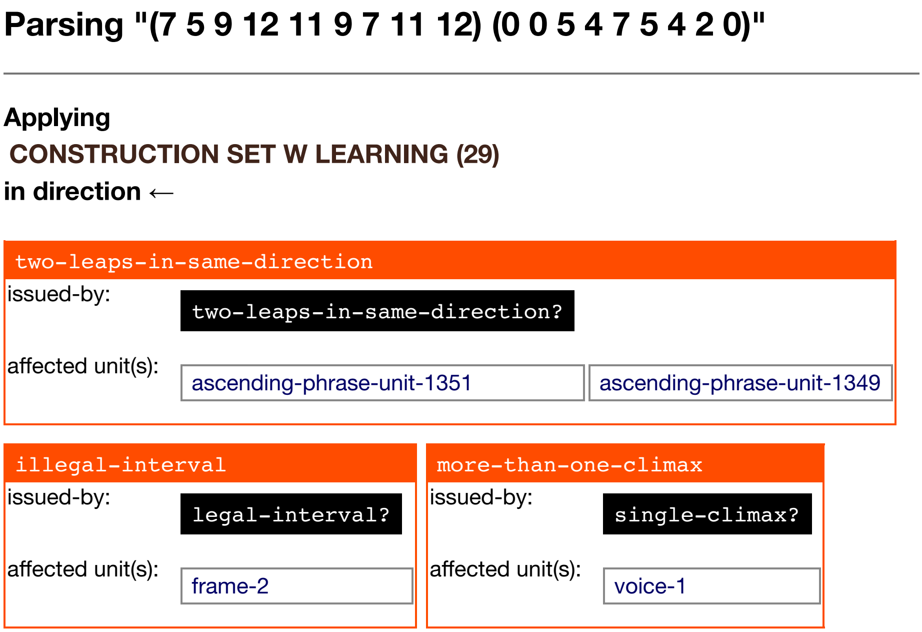

Figure 9. Problems diagnosed for the example melody: two illegal harmonic intervals in frames 2 and 4 and one melodic problem at the end where two skips follow each other in a descending motion (4→2→0).

5.4.1. Harmonic rules

1. **Legal intervals**
   This diagnostic will check all frames (vertical cuts in the piece) to see if each of them is either an octave, perfect fifth, minor/major third or minor/major sixth. Unisons are only allowed in the first bar. There are constructions for each of these intervals as well as for the illegal ones which leave behind a harmonic `interval-category` feature. This feature is a list such as {`third, major-third, consonant`}. All the diagnostic has to check here is the presence of the `consonant` value. In the example melody introduced above, there are two illegal intervals located in frame 2 (perfect fourth) and frame 4 (minor sixth), which is signalled by FCG diagnostics in Figure 8b.

5.4.2. Melodic rules

1. **Big leaps**
   No leaps bigger than a minor sixth (distance 8) are allowed, nor are tritones (distance 6). This diagnostic goes through every phrase unit and checks its melodic interval feature for its distance. If it is bigger than 8, it will diagnose a `leap-bigger-than-minor-sixth` problem. If it is equal to 6, it diagnoses a `tritone` problem.
2. **Exposed tritones**
   A motion where the begin and end note form a tritone is forbidden. To diagnose this, we need to find the largest motion unit (marked with the feature `largest-unit` set to + and then inspect the distance between the first and last

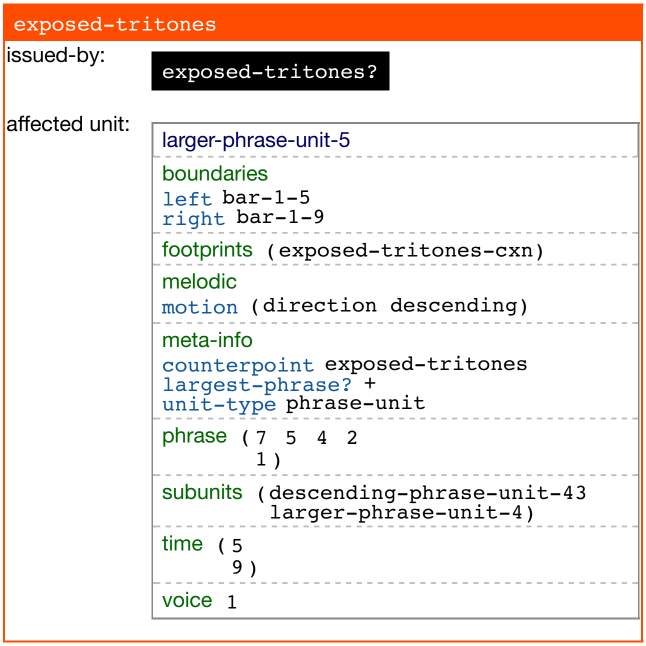

Figure 10. The descending phrase unit contains an exposed tritone, with the difference between start and end note equal to 6 semi-tones. The diagnostic has left a counterpoint feature in the unit's meta info.

note-unit. Instead of manually going through the transient structure, we instantiate a diagnostic construction and check whether it can apply[1]. If it does, the problem is present. We include the diagnostic `exposed-tritones-cxn` here below. Figure 10 shows the problem box in the web interface. Section 5.5 will show in more details how this problem could be repaired.

?top-phrase-unit
meta-info:
  counterpoint: exposed-tritones

←

?top-phrase-unit
meta-info:
  largest-phrase: +
time: {?init-time, ?end-time}
voice: ?voice

?bar-unit-1
pitch
  absolute: ?pitch-1
time: ?init-time
voice: ?voice

>

?bar-unit-2
pitch
  absolute: ++ :exposed-tritones? ?pitch-2
time: ?end-time
voice: ?voice

3. **Climax**
A single climax is preferred, located somewhere in the middle of the counterpoint melody (when it is the highest voice), not at the beginning or the end. This rule requires two steps: first of all there can be only one global climax (marked with

[1]The absolute pitch distance is calculated by means of an expansion operator (indicated by the ++) that is called during the matching process. If no tritone is found, the diagnostic construction cannot apply.

the feature (`global +`). Second, this note should be situated around the middle of the melody. The diagnostic currently defines middle as the length of the melody plus and minus one. For each of these steps a problem can be signalled: `more-than-one-climax` or `single-climax-not-centralized`.

4. **Voice crossing**
   The counterpoint voice should not cross the cantus firmus. As there is a construction that analyses a melody for possible crossings and leaves a feature if there has been one (or multiple), this diagnostic simply has to check the presence of such a feature.

5.4.3. Motion rules

1. **Parallel perfect consonances**
   Thanks to the motion constructions (detecting parallel, similar, contrary or oblique motion in two melodies) and the melodic constructions that mark perfect fifths or octaves, it has become straightforward to signal the presence of a parallel perfect consonance. This diagnostic looks for a parallel motion phrase that contains a perfect consonance feature.
2. **Approaching fifths or octaves**
   A fifth or octave harmonic interval may not be approached by similar motion. This diagnostic again consists of a construction that looks as follows:

$$
\left[\begin{array}{l}
\textbf{?frame-unit} \\ \hline\hline
\text{counterpoint: approaching-5ths-or-8ves}
\end{array}\right]
\leftarrow
\left[\begin{array}{l}
\textbf{?frame-unit} \\ \hline\hline
\text{form:} \\
\quad \text{co-sound: } \{\text{?bar-2, ?bar-4}\} \\
\text{harmonic:} \\
\quad \text{interval:} \\
\qquad \text{category: perfect-fifth } \textit{OR} \text{ octave}
\end{array}\right]
$$

$$
\left[\begin{array}{l}
\textbf{?phrase-unit-1} \\ \hline\hline
\text{phrase: ?phrase-1} \\
\text{melodic:} \\
\quad \text{motion:} \\
\qquad \text{direction: ?same-dir} \\
\text{subunits: } \{\text{?bar-1, ?bar-2}\}
\end{array}\right]
\left[\begin{array}{l}
\textbf{?phrase-unit-2} \\ \hline\hline
\text{phrase: ?phrase-2} \\
\text{melodic:} \\
\quad \text{motion:} \\
\qquad \text{direction: ?same-dir} \\
\text{subunits: } \{\text{?bar-3, ?bar-4}\}
\end{array}\right]
$$

5.5. Suggesting pedagogically sound repairs

Simply indicating whether a counterpoint piece satisfies or breaks particular counterpoint rules without providing clues how to improve a solution, is insufficient in a complex learning domain as music composition. Insightful feedback in the form of "repair strategies" can hence improve learning considerably. Repair strategies formulate ways to solve *individual* problems that were diagnosed by the tutor agent. Solving an individual problem locally, however, can introduce new problems at different points in the melody. Repairing counterpoint violations thus turns into a local search problem that scores repairs and fixes (solutions suggested by a repair strategy).

The main purpose of repairing counterpoint violations is not just to find a solution (or fix) that satisfies all the rules, but one that provides *insight* to the student on how

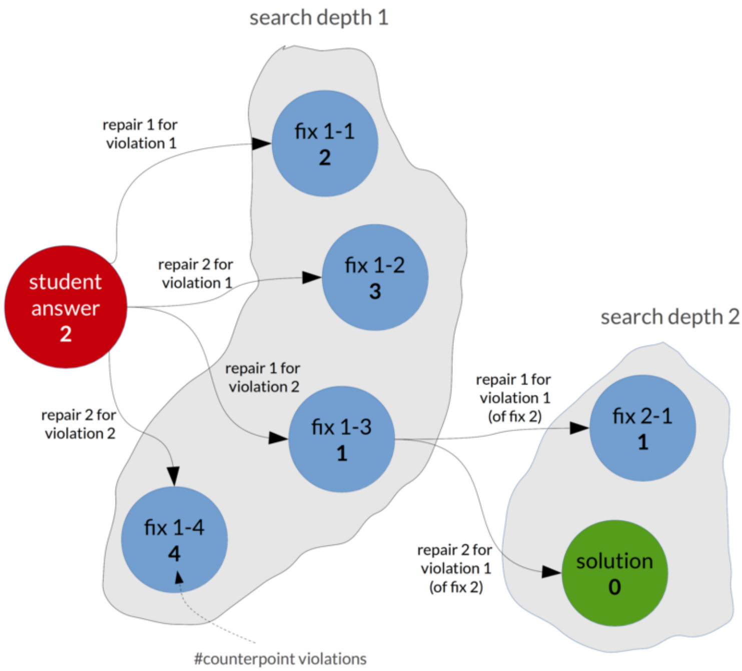

Figure 11. A breadth-first search is applied to find the least number of repairs needed to fix a student's piece. The number of counterpoint rule violations is indicated inside the node. First, all fixes (resulting from a repair) at search depth 1 are explored; next the fixes at depth 2 are explored, ordered by number of violations in ascending order.

to repair his/her mistakes. Therefore, it is important that the tutor provides the "easiest" way to repair its solution. Easy in this sense is the solution where the fewest additional problems were introduced while repairing, i.e. errors that were not made by the student.

Figure 11 illustrates the search process that the tutor runs through to find the best set of fixes. A student piece with two violations, is analysed and all possible fixes (that is, resulting pieces after a repair) are listed together with the remaining number of violations for each. Next, fixes at search depth two are explored and the process is repeated. The tutor follows thus a breath-first search strategy. A simple heuristic of ranking by remaining rule violations is used to choose the order of exploring the nodes at a particular depth. This approach is *guaranteed* to find the *easiest way* to repair the student's piece. Because the original exercise originates from a correct counterpoint with rule violations introduced, the search depth can be limited in practice. For example, when a student has to create a counterpoint voice from scratch, the search depth can be limited as there is no real instructional advantage of proposing many fixes to change a very bad piece into a good one.

It goes beyond the scope of this paper to explain the full details of the tutor's repair strategies. However, we briefly discuss a single problem with its possible repair strategies in what follows. The problem at stake is `exposed tritones` (see above). Figure 12 illustrates the problem and points to three possible repairs:

1. Altering the starting or ending note of the motion phrase.

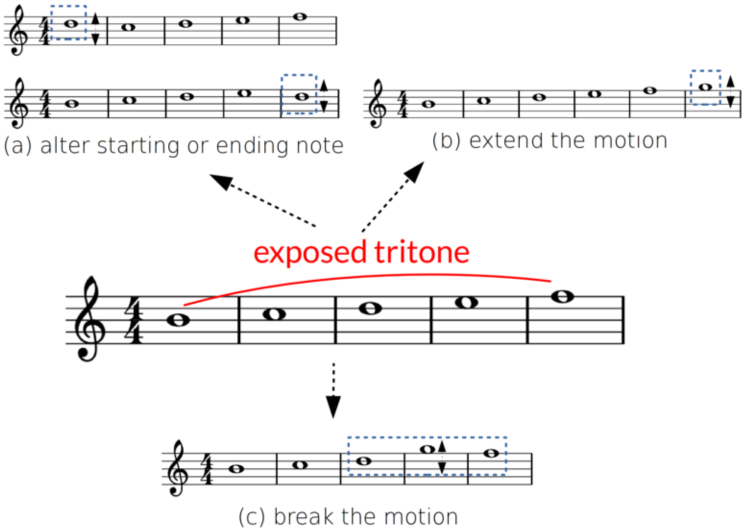

Figure 12. Possible repair strategies for an exposed tritone. The starting or ending note can be changed—with or without breaking the motion (a). The motion can also be extended so the tritone is resolved (b). Another solution is to break the motion within by altering the one-before-last note (c).

2. Extending the phrase so that the tritone is not exposed any more.
3. Breaking the motion somewhere in the middle of the phrase.

These repair strategies are relatively generic and can introduce valid solutions for other problems at the phrase level too (e.g. approaching fifths). They all work on the affected unit(s) for which the problem was diagnosed but also have access to the full transient structure of the musical analysis as they might have to modify other units (corresponding to notes, melodic phrases, harmonic frames or complete voices) in order to solve the counterpoint violation.

For instance, the first strategy to change the initial or ending note of a phrase (given by the problem) yields multiple possible fixes. The exact number is dependent on the range of notes the initial note can be changed into. For now, both the first and the last note of the phrase can be increased or decreased by maximally 3 degrees, resulting in twelve possible fixes for this first repair strategy. To select a valid fix, we rank them based on the number of new problems they introduce into the piece and first pursue the one with the fewest.

5.6. Updating the student model

Finally, the tutor agent has immediate access to the student model so that he can scrutinize the actual state of the agent's construction inventory and learning strategies. It is only then that he can properly align the student model to the real student that is being coached. The student model is aligned after every interaction that the student has with the tutor agent. When the student was successful in the current task, the tutoring strategies will indicate how to update the student agent's constructions that correspond to the task. Indeed, as particular potential rule violation were elicited by the exercise design, the tutor can increase its confidence that the student truly masters these counterpoint rules. Vice

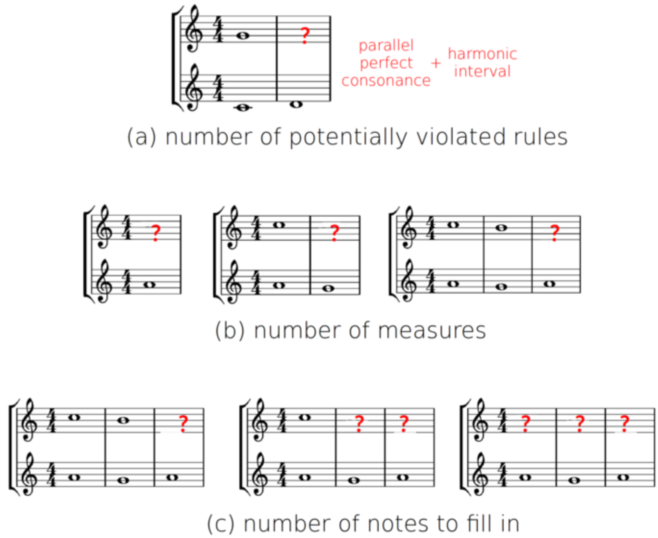

Figure 13. The difficulty level of counterpoint exercises can be controlled in a fine-grained manner. The tutor agent gradually increases the number of potentially violated rules, number of measures and number of notes to fill in, according to the skill level of the student.

versa, when a student's solution is diagnosed with errors, the student model is updated accordingly by lowering the confidence that the student masters the rule.

Also, when there was a failure or mismatch between the goals of the tutor agent and the real student (this can happen specifically when a new difficulty (counterpoint rule) is introduced by the tutor), the student agent proves useful to verify whether the mismatch could have been predicted based on the student model or not. Because the construction inventories have the same architecture in the tutor and the student agent, this symmetry can be used to learn about possible gaps or inconsistencies in the student's grammar.

5.7. Creating counterpoint exercises

The student model informs the tutor about the exercise complexity a pupil can handle and which counterpoint rules (s)he masters. To keep students in flow, fine-grained control over the challenge level is crucial to keep it in balance with the skill level. For this reason, we propose the fine-tuning of three parameters as shown in Figure 13:

(a) the number of measures (longer pieces can introduce more conflicts),
(b) the number of notes to fill in (more degrees of freedom), and
(c) the rules that are potentially violated in the exercise.

The automatic generation of exercises with such a fine-grained particular challenge level, however, has not yet been fully operationalised. In a first stage, we have implemented control measures (a) and (b). In the future, we plan to implement (c) in the following manner. Using a database of correct counterpoint pieces of varying length that is constantly updated, a new piece is generated by altering random notes to create a piece that will probably violate certain rules. Analysis of this piece tells the tutor which rules have been violated exactly. At this point, the tutor can decide to present this piece to the student if it aligns with the particular difficulties the student struggles with (according to the student model), or store it for later use and create a new exercise. Clearly, this approach is far from perfect and needs further refinement.

6. Conclusions

We can now speak of at least two generations of ITS research, and we are currently at the dawn of a third one, which will probably be much more revolutionary. The first generation spans roughly from 1970 until 1990, a period of thirty years in which more than 40 systems were released. This early generation was powered by the booming of Artificial Intelligence, a field that was seeking applications for its technologies. The second generation, ranging from roughly 1990 until today has formulated the scientific foundations of the field [58]. Also implementations of real systems in schools realized, thanks to new spread of digital technologies in traditional education. The third generation massively scales the potential of Computer-Assisted Instruction to online video-based courses that are freely accessible to thousands of students: the so-called MOOCs. Although the first hype of MOOCs in the years 2012-2013 had set high hopes onto the disrupting nature of these courses on the traditional education landscape, they failed to accomplish them.

One reason for their failure that this chapter has put forward is the lack of individual tutoring they offer (when used in a distance education setting), which is mainly due to the absence of a student model and tutoring strategies. By building a bridge to earlier techniques found in Intelligent Tutoring Systems and enhancing these by making use of a truly predictive student model in the form of an active autonomous agent that simulates the actual learner.

In this chapter, we have outlined a tutoring system based on the theory of flow. Flow theory describes the experience of intrinsically motivated people when they are deeply engaged in an activity. It is believed that learning is optimized when students are in a state of flow. This situation happens when the challenge level is kept in balance with the skill level of the student to avoid boredom and anxiety. For this reason, an architecture based on active tutor and student agents is proposed. As the student agent carefully track the pupil's skill level and can simulate its answers, appropriate exercises can be presented to the student that keeps her/him in a state of flow.

The proposed methodology has been tested in the domain of music, more precisely the Study of Counterpoint, still today an effective instructional tool to teach students the craft of polyphonic music composition. An operational grammar of music and counterpoint as well as a tutor and student agent has been implemented in Fluid Construction Grammar (FCG). A core component of the virtual tutor is the in-depth musical analysis of the student's piece and the possibility to suggest repairs that correct the mistakes. This way, the student does not only get feedback on the correctness of an answer (right/wrong), but also insight into his/her mistakes and how to solve them, a mechanism central to any learning process. Though the creation of counterpoint exercises that elicit specific counterpoint violations is still suboptimal, the proposed work opens up the way to a stand-alone online counterpoint tutor.

References

[1] L. Steels, "The coming of moocs," in *Learning about music in the age of MOOCS* (L. Steels, ed.), Amsterdam: IOS Press, 2015.

[2] B. F. Skinner, *Cumulative record*. New York: Appleton Century Crofts, 1961.

[3] B. Bloom, "The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring," *Educational researcher*, vol. 13, pp. 3–16, 1984.

[4] G. Ritzer, "The rise of the prosuming machines." Blog article, March 2014.

[5] K. Colvin, J. Champaign, A. Liu, Q. Zhou, C. Fredericks, and D. Pritchard, "Learning in an introductory physics mooc: All cohorts learn equally, including an on-campus class," *The International Review of Research in Open and Distance Learning*, vol. 15, no. 4, 2014.

[6] D. Koller, "Death knell for the lecture: Technology as a passport to personalized education," *New York Times*, 5 December 2011.

[7] F. G. Martin, "Will massive open online courses change how we teach?," *Commun. ACM*, vol. 55, pp. 26–28, Aug. 2012.

[8] J. Kay, P. Reimann, E. Diebold, and B. Kummerfeld, "Moocs: So many learners, so much potential ...," *IEEE Intelligent Systems*, vol. 28, no. 3, pp. 70–77, 2013.

[9] S. Kahn, "Let's use video to reinvent education." TED Talk, March 2011.

[10] J. Bowman, *Online Learning in Music: Foundations, Frameworks, and Practices*. Oxford University Press, 2014.

[11] A. Agarwal, "Why massive open online courses (still) matter." TED Talk, June 2013.

[12] J. A. Self, "Bypassing the intractable problem of student modeling," in *Intelligent Tutoring Systems: At the crossroads of Artificial Intelligence and Education* (C. Frasson and G. Gauthier, eds.), Norwood, NJ: Ablex, 1988.

[13] N. A. Crowder, "On the difference between linear and intrinsic programing," in *Programs, teachers, and machines* (A. G. Grazia and D. A. Sohn, eds.), pp. 77–85, New York, NY: Bantam Books, 1964.

[14] L. Uhr, "Teaching machine programs that generate problems as a function of interaction with students," in *Proceedings of the 24th National Conference,*, pp. 125–134, 1969.

[15] J. Carbonell, *Mixed-initiative man-computer instructional dialogue*. PhD thesis, MIT, Cambridge, MA, 1970.

[16] J. Carbonell, "Ai in cai: An artificial intelligence approach to computer aided instruction," *IEEE Transactions on Man-Machine Systems*, vol. 11, pp. 190–202, 1970.

[17] W. J. Clancey, *Transfer of Rule-based Expertise through a Tutorial Dialogue*. PhD thesis, Stanford University, Stanford, CA, 1979.

[18] J. A. Self, "Student models in computer aided instruction," *International journal of man-machine studies*, vol. 6, pp. 261–276, 1974.

[19] J. A. Self, "Concept teaching," *Artificial Intelligence*, vol. 9, pp. 197–221, 1977.

[20] J. A. Self, "A perspective on intelligent computer-aided instruction," *Journal of Computer Assisted Learning*, vol. 1, pp. 159–166, 1985.

[21] J. Anderson, A. Corbett, K. Koedinger, and R. Pelletier, "Cognitive tutors: Lessons learned," *The journal of the learning sciences*, vol. 4, no. 2, pp. 167–207, 1995.

[22] A. Lesgold, S. Lajoie, M. Bunzo, and G. Eggan, "Sherlock: A coached practice environment for an electronics troubleshooting job," *Computer Assisted Instruction and Intelligent Tutoring Systems*, pp. 201–238, 1992.

[23] W. J. Clancey, *Knowledge-based Tutoring: The Guidon Program.* Cambridge, MA: MIT Press, 1987.

[24] B. Park Woolf, *Building Intelligent Interactive Tutors.* Burlington, MA: Morgan Kaufmann, 2008.

[25] R. Cook and J. Kay, *The justified user model: a viewable, explained user model.* Basser Department of Computer Science, University of Sydney, 1994.

[26] K. VanLehn, "Student Modeling," in *Foundations of Intelligent Tutoring Systems* (M. Polson and J. Richardson, eds.), pp. 55–78, Hillsdale, NJ: Erlbaum, 1988.

[27] M. Polson and J. Richardson, *Foundations of Intelligent Tutoring Systems.* Interacting with Computers Series, Taylor & Francis, 2013.

[28] K. VanLehn, "Intelligent tutoring systems for continuous, embedded assessment," in *The future of assessment: Shaping teaching and learning* (C. Dwyer, ed.), Lawrence Erlbaum Associates, 2007.

[29] L. Razzaq, J. Patvarczki, S. Almeida, M. Vartak, M. Feng, N. Heffernan, and K. Koedinger, "The assistment builder: Supporting the life cycle of tutoring system content creation," *Learning Technologies, IEEE Transactions on*, vol. 2, no. 2, pp. 157–166, 2009.

[30] L. Razzaq and N. Heffernan, "Open content authoring tools," *Advances in Intelligent Tutoring Systems*, pp. 407–420, 2010.

[31] T. Turner, M. Macasek, G. Nuzzo-Jones, N. Heffernan, and K. Koedinger, "The assistment builder: A rapid development tool for its," in *Proceedings of the 12th Annual Conference on Artificial Intelligence in Education*, pp. 929–931, 2005.

[32] J. Beck, K. Chang, J. Mostow, and A. Corbett, "Does help help? introducing the bayesian evaluation and assessment methodology," in *Intelligent Tutoring Systems*, pp. 383–394, Springer, 2008.

[33] L. Razzaq and N. Heffernan, "To tutor or not to tutor: That is the question," in *Proceedings of the Conference on Artificial Intelligence in Education*, pp. 457–464, 2009.

[34] J. Bourdeau and M. Grandbastien, "Modeling tutoring knowledge," in *Advances in Intelligent Tutoring Systems* (R. Nkambou, J. Bourdeau, and R. Mizoguchi, eds.), pp. 123–143, Springer, 2010.

[35] I. Arroyo, D. G. Cooper, W. Burleson, B. P. Woolf, K. Muldner, and R. Christopherson, "Emotion sensors go to school," in *Proceeding of the 2009 conference on Artificial Intelligence in Education, July 6th-10th, Brighton, UK, IOS Press*, pp. 17–24, 2009.

[36] J. Walonoski and N. Heffernan, "Detection and analysis of off-task gaming behavior in intelligent tutoring systems," in *Intelligent Tutoring Systems*, pp. 382–391, Springer, 2006.

[37] B. Woolf, I. Arroyo, D. Cooper, W. Burleson, and K. Muldner, "Affective tutors: Automatic detection of and response to student emotion," *Advances in Intelligent Tutoring Systems*, pp. 207–227, 2010.

[38] A. Ferreira, J. D. Moore, and C. Mellish, "A Study of Feedback Strategies in Foreign Language Classrooms and Tutorials with Implications for Intelligent Computer-Assisted Language Learning Systems," *International Journal of Artificial Intelligence in Education*, vol. 17, pp. 389–422, 2007.

[39] R. Lyster and L. Ranta, "Corrective Feedback and Learner Uptake," *Studies in Second Language Acquisition*, vol. 19, pp. 37–66, Mar. 1997.

[40] I. Panova and R. Lyster, "Patterns of corrective feedback and uptake in an adult ESL classroom," *Tesol Quarterly*, vol. 36, no. 4, pp. 573–595, 2002.

[41] A. C. Graesser, P. Chipman, B. C. Haynes, and A. Olney, "AutoTutor: an intelligent tutoring system with mixed-initiative dialogue," *IEEE Transactions on Education*, vol. 48, Nov. 2005.

[42] R. Nkambou, J. Bourdeau, and P. V., "Building Intelligent Tutoring Systems: An Overview," in *Advances in Intelligent Tutoring Systems* (R. Nkambou, J. Bourdeau, and R. Mizoguchi, eds.), Studies in Computational Intelligence, ch. 18, pp. 361–375, Berlin: Springer, 2010.

[43] S. Sun, M. Joy, and N. Griffiths, "The use of learning objects and learning styles in a multi-agent education system," *Journal of Interactive Learning Research*, vol. 18, no. 3, pp. 381–398, 2007.

[44] M. A. Razek, C. Frasson, and M. Kaltenbach, "Toward more cooperative intelligent distance learning environments," *Software Agents Cooperation Human Activity*, 2002.

[45] M. D. Beer and J. Whatley, "A multi-agent architecture to support synchronous collaborative learning in an international environment," in *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*, pp. 505–506, ACM, 2002.

[46] Y. Shang, H. Shi, and S.-S. Chen, "An intelligent distributed environment for active learning," *Journal on Educational Resources in Computing (JERIC)*, vol. 1, no. 2es, p. 4, 2001.

[47] M. Boicu, G. Tecuci, B. Stanescu, D. Marcu, M. Barbulescu, and C. Boicu, "Design principles for learning agents," in *Proceedings of AAAI-2004 Workshop on Intelligent Agent Architectures: Combining the Strengths of Software Engineering and Cognitive Systems*, pp. 26–33, 2004.

[48] K. Beuls, *Towards an agent-based tutoring system for Spanish verb conjugation*. PhD thesis, Vrije Universiteit Brussel, November 2013.

[49] J. J. Fux, *The study of counterpoint from Johann Joseph Fux's Gradus ad Parnassum*. No. 277, WW Norton & Company, 1965.

[50] P. Hindemith, *The Craft of Musical Composition*, vol. 2. Schott & Co Ltd, 1984.

[51] A. E. Yilmaz and Z. Telatar, "Note-against-note two-voice counterpoint by means of fuzzy logic," *Knowledge-Based Systems*, vol. 23, no. 3, pp. 256–266, 2010.

[52] T. Anders and E. R. Miranda, "Constraint programming systems for modeling music theories and composition," *ACM Computing Surveys (CSUR)*, vol. 43, no. 4, p. 30, 2011.

[53] D. Herremans and K. Sörensen, "Composing fifth species counterpoint music with a variable neighborhood search algorithm," *Expert systems with applications*, vol. 40, no. 16, pp. 6427–6437, 2013.

[54] M. Csikszentmihalyi, *Beyond bordedom and anxiety: experiencing flow in work and play*. Cambridge University Press, 1978.

[55] L. Steels, "The architecture of flow," in *A Learning Zone of One's Own* (M. Tokoro and L. Steels, eds.), pp. 137–149, IOS Press, 2004.

[56] L. Steels, *Design patterns in fluid construction grammar*, vol. 11. John Benjamins Publishing, 2011.

[57] L. Steels, "Basics of fluid construction grammar," *Constructions and Frames*, 2015.

[58] J. A. Self, "Theoretical foundations for intelligent tutoring systems," *Journal of Artificial Intelligence in Education*, vol. 1, no. 4, pp. 3–14, 1990.

This page intentionally left blank

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-145

Chapter 10. Collaborative Peer Assessment using PeerLearn

Ismel Brito [a] Patricia Gutierrez [a] Katina Hazelden [a] Dave de Jonge [a] Lisette Lemus [a] Nardine Osman [a] Bruno Rosell [a] Carles Sierra [a] and Carme Roig [b]
[a] *IIIA-CSIC, Bellaterra, Catalonia, Spain*
[b] *IES Torras i Bages, L'Hospitalet, Catalonia, Spain*

AbstractIn this chapter we introduce the PeerLearn methodology and its associated tools. We base the design of pedagogical workflows for students on the definition of rubrics (using PeerAssess) as the starting element that drives the creation of lesson plans (using LessonEditor). These plans run over our web platform (PeerFlow). Students can evaluate one another following given rubrics and teachers can accept (or not) marks produced by a collaborative assessment tool (COMAS). Experimental results show that PeerLearn provide students with a highly satisfying new pedagogical experience and increased learning outcomes.

1. Introduction

There are a number of available tools that support teachers in the management of lesson plans on the web. However, none of them is task-centered and support any form of lesson plan's 'execution' over the web. PlanBoard[1] and PlanbookEdu[2] deal with issues such as lesson planning, standards setting, assessment management, etc. that help teachers in their scheduling and management of resources. CorePlanner[3] on top of that allows teachers to set the objectives of classes following national standards like the Core Standards.[4] There is a large number of repositories of lesson plans that can be consulted freely.

All mentioned tools understand a lesson plan as a document that can be shared among teachers. They do not provide any IT environment where these lesson plans become 'executable'. Against this background, our goal in this paper is to introduce PeerLearn, a set of tools for the design and, most importantly, the *execution* of lesson plans.

Thus, In this paper we propose a methodology and associated tools suite that covers the educational process that goes from the preparation of evaluation rubrics to the actual assessment of students (see Figure 1).

The methodology consists of the following steps:

- **Define a rubric**. In modern pedagogy [9] the definition of a rubric is the first step in the preparation of teaching materials. How students will be evaluated has to

[1]https://www.planboardapp.com/
[2]http://planbookedu.com/
[3]http://coreplanner.com/
[4]http://www.corestandards.org/

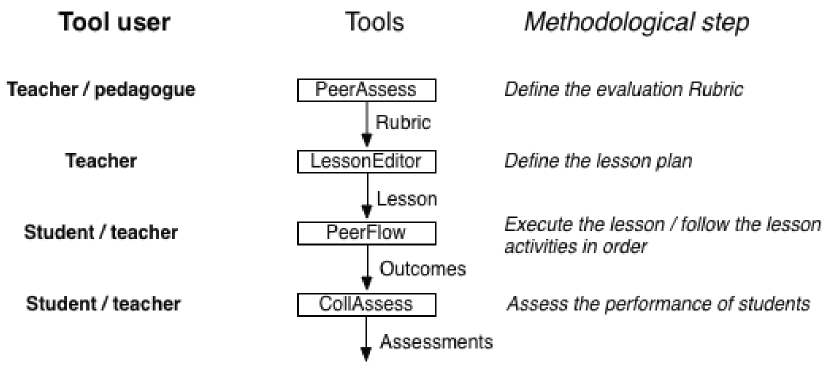

Figure 1. Methodology and PeerLearn tools for lesson plan management.

drive the relevance of the learning materials and activities to be defined by teachers. The tool PeerAssess, explained in detail in Section 2.6, helps pedagogues in the management of a database of rubric criteria. Pedagogues, or teachers, may create new evaluation criteria, adapt generic criteria or group criteria as a rubric.

- **Define a lesson plan**. In order to achieve the pedagogical objectives associated with a rubric teachers must lay down a sequence, or flow, of activities. These activities can be individual or in groups. The performance of an activity may need the consumption of digital resources, e.g. files containing explanations, exercises, etc., and may require that students produce outcomes like presentations, audio or video recordings or documents. We have developed a tool called FlowEditor that allows for the definition of sophisticated lesson plans. The tool is explained in detail in Section 3.3.
- **Do the lesson**. Once a lesson plan is defined, students have to follow the activities of the lesson plan in order to achieve the desired competencies and produce the outcomes that would allow for their evaluation. The tool PeerFlow allows for the assignment of lesson plans to groups of students, imposes restrictions on the activity flow of students so that they perform them in the right order, grants access to the digital resources established in the lesson, and keeps a database of the files produced by the students. PeerFlow facilitates the structured interactions that the lesson plan includes, like group communication, teacher check points, and so on. The tool is explained in detail in Section 4.5.
- **Assess students**. Finally, we have developed an assessment tool called COMAS that allows the mutual assessment of students and uses a measure of the teacher's trust on the evaluation skills of students to assign marks to students according to the rubric, minimising the number of assessments that teachers have to make. This is specially relevant when the number of students is very large as in the recent MOOCs movement. Section 5.3 provides the details of this tool.

2. PeerAssess

In this Section we describe the tool PeerAssess. PeeAssess is an online database platform that helps teachers in the management of groups, lesson plans and evaluation criteria (rubrics), and makes defining subsequent lesson plans easy. In the following we enumerate the functionalities embedded in PeerAssess and provide examples of the pedagogical materials that can be managed through this platform.

2.1. Log in

Teachers log in to PeerAssess through the Welcome page (Figure 2). Once logged in, they have access to their personalized information and can manage their groups, their lesson plans and assessment materials (rubrics).

2.2. Manage Groups

As teachers can teach in one or more groups, Group management (add, edit, remove, list groups) can be performed choosing the "Groups" menu option (Figure 3).

2.3. Manage Evaluation Criteria

Formative assessment plays a very important role during the learning process and students should be frequently assessed. Assessment is usually performed upon different skills based on different assessment criteria. For instance, in an English course, oral skills may be evaluated considering one or more of the following criteria: fluency, pronunciation, clarity, speed, etc. The definition of the assessment criteria is important because it provides a clear view of the desired outcome and of how an assignment will be graded (something which the student often considers to be subjective). The setting of clear assessment criteria simplifies the marking process and helps ensure that marking is a thoughtful exercise mapped to specific desired learning outcomes. In addition, the definition of the assessment criteria plays an important role in the preparation of the learning materials and activities that will be used in the lessons.

In PeerAssess teachers are able to create personalized assessment criteria that may be later on used in different assignments. Assessment criteria consist of a name and a

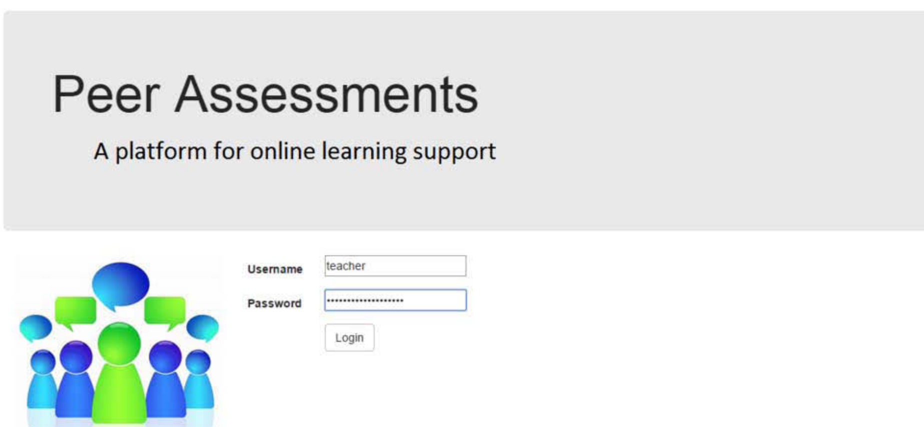

Figure 2. Log in Page

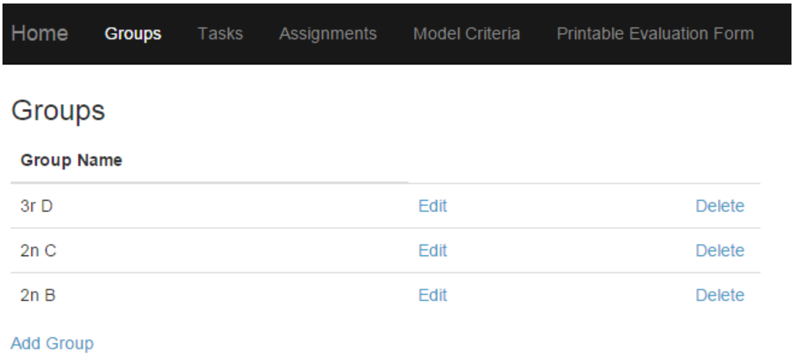

Figure 3. Group Management

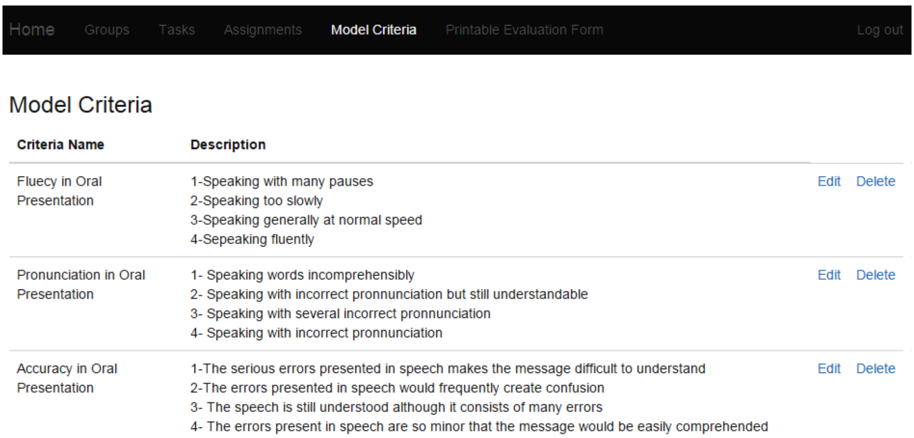

Figure 4. Evaluation Criteria Management

graded description that specifies the requirements to be met by students' performance, from completely met to "not satisfactorily met" (usually 4 descriptors, see examples in Figure 4). Teachers can manage the evaluation criteria (add, edit, remove, list) choosing the "Model Criteria" option from the menu.

2.4. Manage Tasks

Teachers are able to create different lesson plans. Lesson plans consist of a number of tasks conducent to a learning outcome. A lesson is defined by its name and a description of the outcomes that students will have to deliver (see examples in Figure 5). Teachers can manage their lessons (add, edit, remove, list tasks) choosing the "Tasks" menu option.

2.5. Manage Assignments

Once the teacher has decided the final outcome, the assessment criteria and the different tasks — that is, once the teacher has a lesson plan — specific lesson plans can be cus-

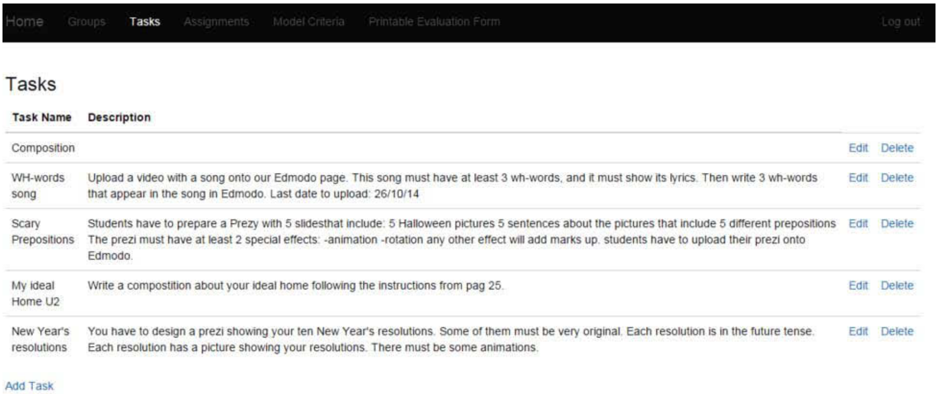

Figure 5. Task Management

tomized for different groups. An assignment is defined by a name, a submission date, an outcome (which describes the activity), the group that will be assigned to, and a set of assessment criteria (a rubric) that measures the student's performance. The set of assessment criteria defines a *rubric*. Rubrics are of great importance to ensure the consistency of assessments along different dates of marking and among different markers. Rubrics guide students in the understanding of the assignment, making explicit what is required, what qualities and skills are looked for, the weight of these skills in the marking process and the essential elements that shouldn't be missed.

Teachers can manage their assignments (add, edit, remove, list) choosing the "Assignments" menu option. To make the rubric creation more user friendly and avoid the need to repeat the same evaluation criteria from one assignment to the next, we provide the functionality of selecting an already created evaluation criteria (see functionality **Manage Evaluation Criteria**) and loading the stored information from that criteria into the form. Teachers can then add the selected criteria directly into the rubric, or modify it and add it into the rubric. In Figure 6 an example of assignment creation and criteria selection for the rubric using a pre-defined evaluation criteria is shown.

2.6. Print or Download Evaluation Form

Once an assignment has been defined, teachers can generate an evaluation form for its assessment choosing the "Printable Evaluation Form" menu option (Figure 7) and filling the group and assignment filters. The way in which assignments will be delivered during the course and the way in which evaluations will be conducted (in-person or online assignments, group submissions or individual, peer-to-peer or teacher evaluations, etc) will depend on the lesson plan designed by the teacher. In any case, this tool helps the tutor providing an evaluation template for assignments, whether in a printable format, using the "Show Evaluation Form" option, or for online evaluations, using the "Download rubric to xml" option. In the latter case, the xml file generated will be used in the Peerflow tool (see next Sections) when defining the rubric in the lesson plan.

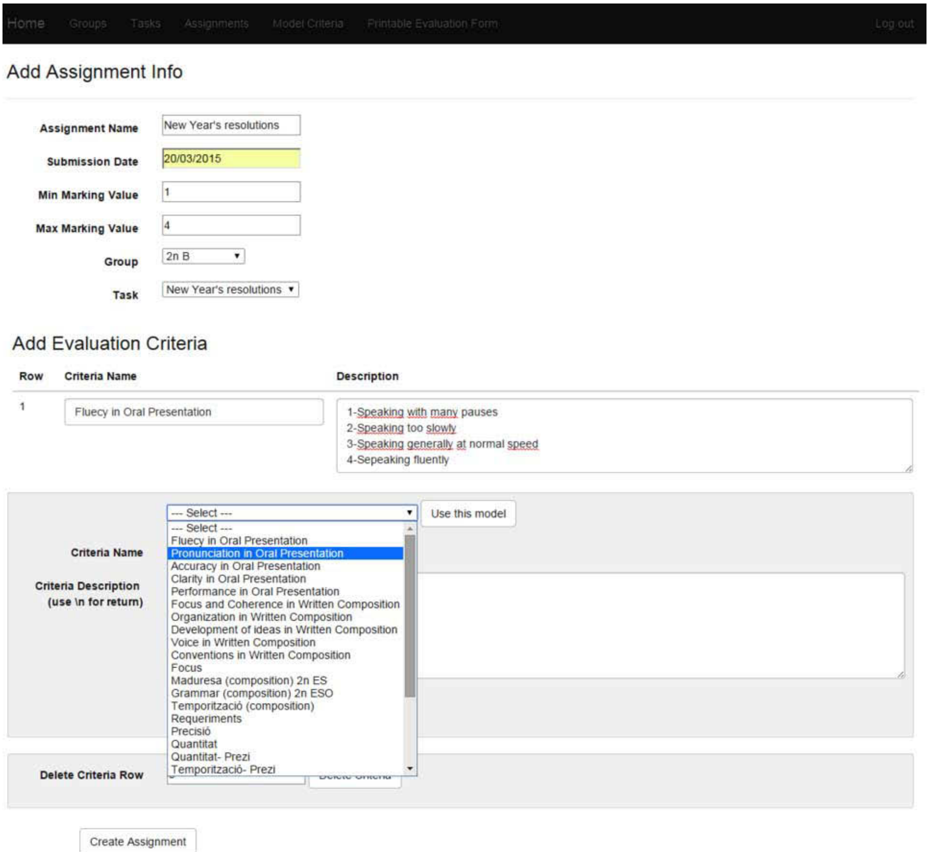

(a) Add Assignment

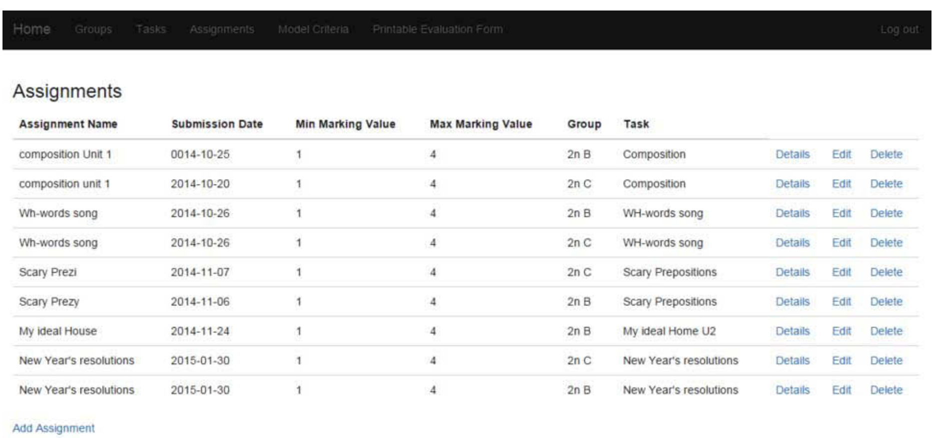

(b) List Assignments

Figure 6. Assignment Management

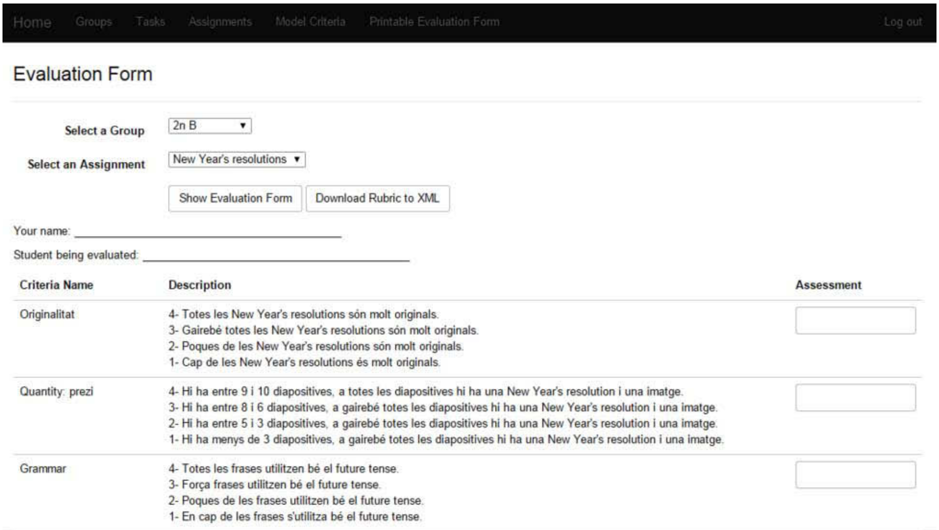

Figure 7. Evaluation Form

3. Lesson Plan Editor

This section is intended to give assistance to people creating workflows using the web-based PeerFlow Editor. We envision a workflow in PeerFlow as a recipe with six key ingredients: activities, roles, resources, applications, timeouts and files. Broadly speaking, we define a workflow as an ordered set of activities that should be carried out by its participants, playing a certain role. Resources, applications, timeouts, and files allow us to further enrich the definition of the actions that participants can perform within activities. In a user-friendly manner, the editor enables people to drag-and-drop and link different graphical elements to create workflows.

In the following subsections we introduce the six key elements aforementioned through an example. Furthermore, we provide guidance for the design of a workflow using the PeerFlow Editor. We use the example of a class room environment where a teacher and a number of students engage in a set of educational activities to clarify the concepts. Finally, we introduce the New Year's resolution workflow, which have been designed and tested in the PeerFlow platform.

Throughout this section we will refer to the person designing a workflow as 'the workflow designer' or simply 'the designer'.

3.1. *Ingredients of a PeerFlow Workflow*

A workflow is an interaction model that determines the way its participants perform a set of activities in a predefined order. Designing a workflow can be a very challenging task without the appropriate supporting tools. Obviously, a workflow designer would be frustrated if a supporting tool requires him or her to know many technical details. This is in fact the main issue we face in the context of the PRAISE project, where we want music teachers with average computer skills to be able to design lesson plans. The PeerFlow Editor helps people to create workflows in an easy-to-use and user-friendly environment.

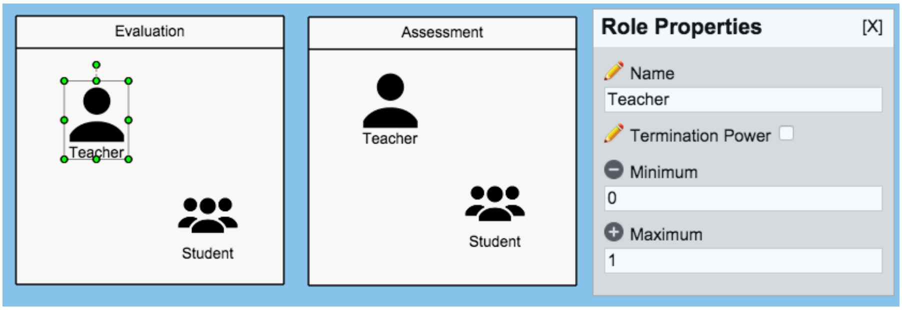

Figure 8. Evaluation and Assessment Activities

In the following subsections we introduce and provide examples of the six basic elements of PeerFlow workflows: activities, roles, resources, applications, timeouts and files. We also provide an example of the use of group activities. A group activity is an activity in which the participants are split into groups of a desired size.

3.1.1. Activities and Roles

One of the central concepts in PeerFlow is the concept of a *role*. Every participant in a workflow is required to adopt one of the roles that have been defined by the designer of the workflow. The role of a participant determines what he or she is allowed to do. Two obvious examples of roles in a teaching environment would be the roles ‘teacher’ and ‘student’.

Another central concept in PeerFlow is the concept of an *activity*. A workflow is in fact a sequence of activities, and each activity defines what actions can be performed by the participants involved in it and what resources are available to them.

For instance, let us consider a simple class room scenario. For this example we can create a workflow with two activities, named Evaluation and Assessment. In the former activity the teacher assigns students an assignment and each student returns his or her answers to the assignment. In the latter activity the teacher assesses each student’s assignment. Both activities involve the teacher as well as the students, so we need to specify for both activities that the teacher as well as the students can enter. Figure 8 displays the aforementioned activities and the data the designer needs to fill in for each role in the activity. Besides the role name, the designer should define whether this role is permitted to terminate the activity and the minimum and maximum number of participants that can adopt this role.

3.1.2. Resources, Applications and Timeouts

Activities may also include resources, applications and a timeout. A *resource* is a document that participants in the workflow may consult within the activity. In the example above, the designer could decide to give the students access to an online dictionary during the Assessment activity. This can be achieved by adding a resource to the Evaluation activity that links that dictionary. As can be seen in Figure 9, each included resource is determined by a name, a description and an the Internet address where the resource is located.

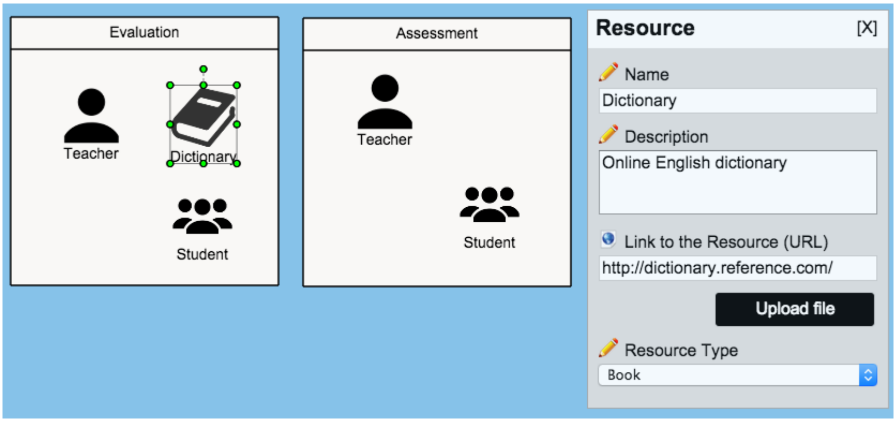

Figure 9. Adding a dictionary as resource to the Evaluation activity.

To improve the attractiveness and readability of workflows, the editor permits the designer to choose among 5 different icons to represent the type of resource: book, video, image, audio and document. Changing the icon does not change the functionality of the resource. However, choosing the appropriate icon helps the designer to see what kind of resource it represents. By default, each new resource added to a workflow is represented by the book icon, which can then be changed as desired.

An *application* in PeerFlow is an external service running online that, when consulted, allows users to perform specific tasks. One can see an application as an online software tool that provides extra functionality that is not part of the PeerFlow framework. Unlike resources, which are passive objects, the participants may need to interact with an application in the activity according to some protocol defined by the application itself. In the music learning domain (and hence in the context of the PRAISE project) the use of applications is of paramount importance. For instance, there are software tools for automatically assessing the performance of musical recordings. It is here where the use of applications is vital to connect external tools with the workflows defined in PeerFlow. In this way, the workflow designer enriches its workflows by means of applications, which provide extra functionalities to workflows. In the example considered, one can imagine an application that supports teachers in automatically assessing students' exams. For the sake of simplification, however, we decided not to include any application in this example. We will return to applications when we introduce the new year resolution example in subsection 3.3.

Before continuing, let us review the approach we proposed in the example above. We mentioned the need of having two activities: Evaluation and Assessment. The Evaluation activity covers two main actions: the teacher handing out the assignment to the students and the students returning their finished assignments. As is common in teaching, the teacher may want his students to have a deadline for returning the assignment. In order to address this challenge and give clarity to our example we now propose to split the first activity into two new ones: Assignment and Examination. A *timeout*, one of the other ingredients of the PeerFlow recipe, is what we need to face the teacher's new requirement. We have redrawn the activities of our example in Figure 10. Notice that the

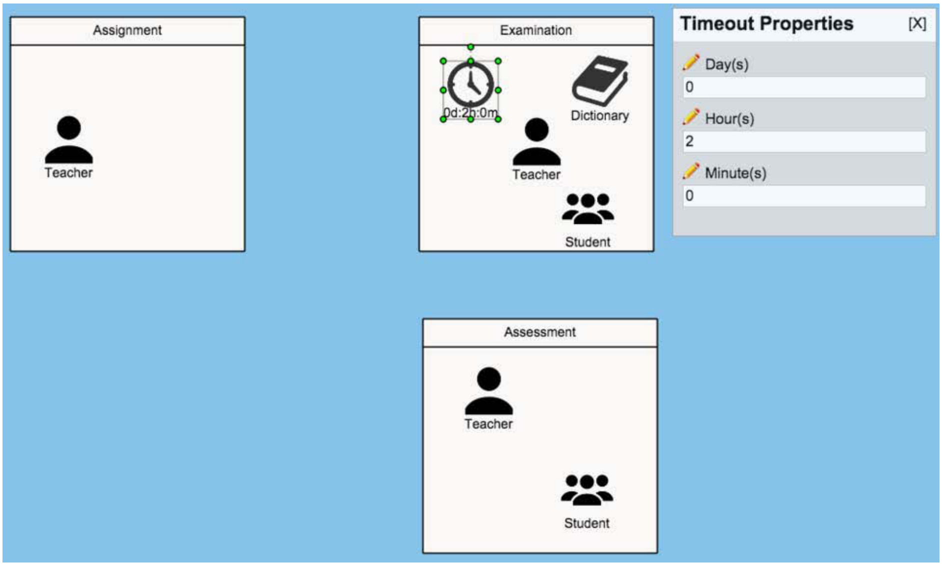

Figure 10. Adding a timeout to the Examination activity.

workflow designer should define a timeout as the number of days, hours and minutes the activity lasts. In this case, we state the students have 2 hours to submit the assignment.

3.1.3. File sharing

A *file* is a concept similar to a resource. Just like a resource, it represents an online document that is available to the participants. The difference however is that a resource is an immutable document, while a file can be created or adapted at runtime by the participants. There are two kinds of files: intra-files and inter-files. An intra-file is generated and consulted by users inside a single activity, while inter-files can be shared across multiple activities. Each inter-file must be linked to two roles belonging to two different activities. Those links are established by arrows in the editor. In Figure 11, we show how two inter-files can be used to define the process of a teacher handing out an exam to its students, followed by the students handing their answers back in to the teacher. The arrow from the teacher to the exam file indicates that the teacher must upload this file in the Assignment activity. The arrow from the exam file to the students, on the other hand, indicates that the students in the Examination activity can download the exam file. Once the students finish their exams, they should be able to send back their answers to the teacher. In order to make this possible we have added an arrow from the students to the answers file.

It might seem that our workflow is complete, but if we go back to the definition of the workflow we gave in the first paragraph of this section, we should note that something is lacking. We have not yet defined the order in which the activities take place. Of course, in this example the order of the activities could be inferred from the directions of the arrows between the inter-files and the roles. However, not every workflow may contain inter-files. For this reason we also need to explicitly indicate how the users can move from activity to activity.

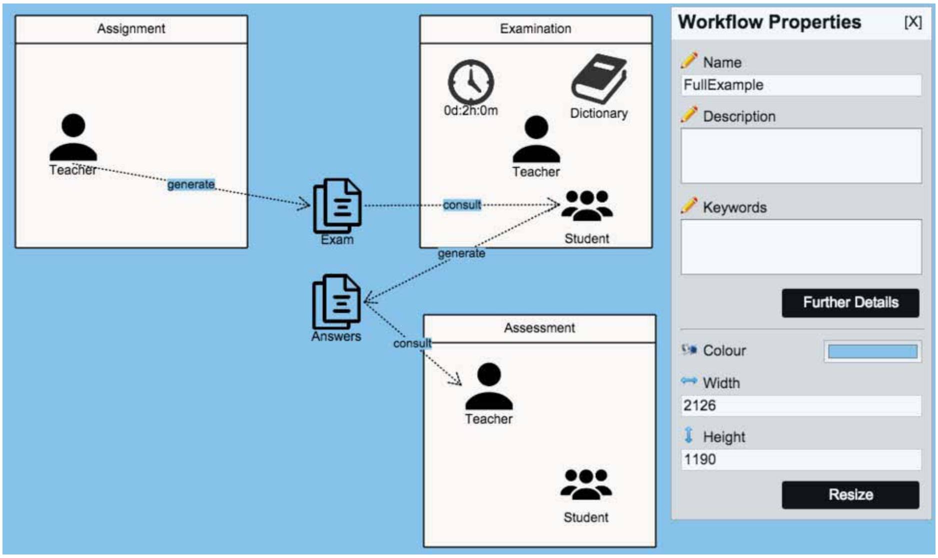

Figure 11. Adding Exam and Answers as files that are exchanged between teacher and students.

The order between activities can be established by drawing arrows between roles in different activities. In this way, the designer can determine the order in which the users of a specific role can enter the different activities. In Figure 12 we have added such arrows to our example. The new workflow indicates that the process starts in the Assignment activity in which only a teacher is involved. Once this activity is finished, the teacher is then permitted to move to the Examination activity and eventually to the Assessment activity. The students on the other hand, start in the Examination activity and when they have finished their exams, they can move to the Assessment activity.

3.1.4. Group activities

Another common requirement in teaching environments is to split the class into smaller groups of students and have the students working together, as a group, on a certain assignment. In Figure 13 we have changed the previous example to set the Examination activity as a group activity where the ‘groups’ are of size 1. After all, taking an exam is indeed an individual activity.

3.2. The PeerFlow Editor

In this section we introduce the PeerFlow Editor, as illustrated by Figure 14. The editor window is split into five working areas: the menu bar, the toolbars that appear at the top and left side of the window, the main screen or edition area, and the properties panel. The menu bar contains a set of actions for creating, saving, importing and publishing workflows into the PeerFlow site, whereas the toolbar at the top of the main screen includes the shortcuts to some of the actions appearing in the menu bar, as highlighted by Figure 15. Of special interest is the fourth shortcut in this toolbar that allows users to draw arrows among the basic elements. We will discuss workflow arrows below.

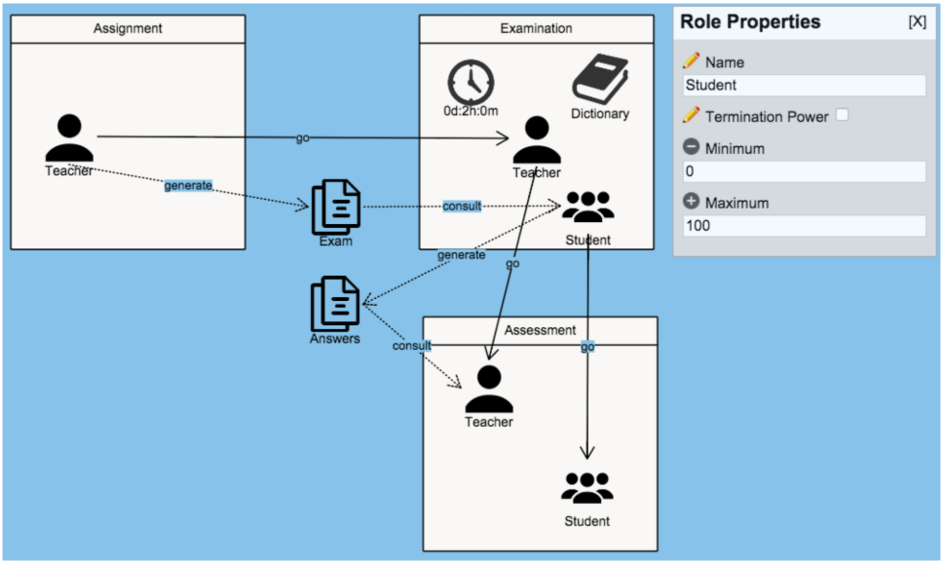

Figure 12. Setting the order in which the participants playing a certain role may enter the activities.

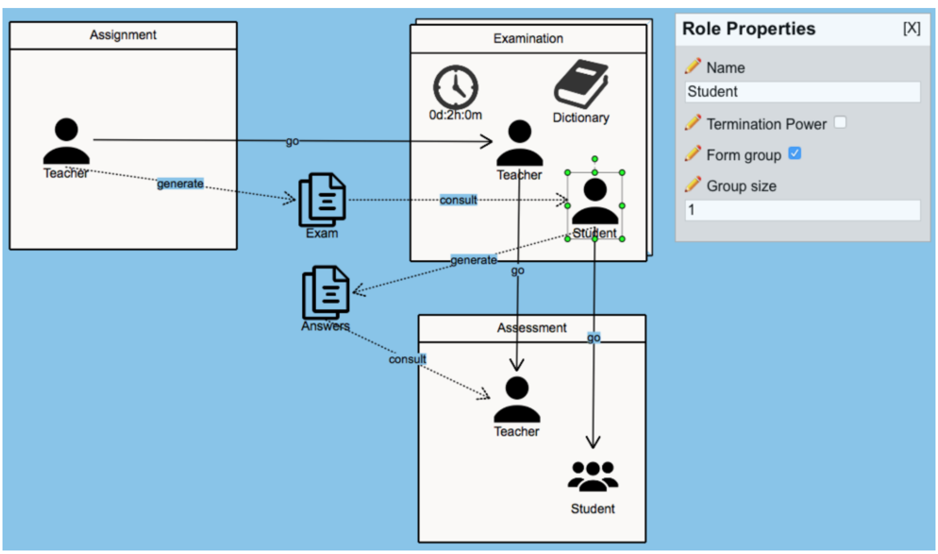

Figure 13. Changing Examination as an activity in which each student takes the exam alone.

The six necessary elements for creating a workflow are shown in the left-side toolbar of 14, and highlighted by Figure 16. Those are: activities, roles, resources, applications, files and timeouts. The main screen (14) is the area where the designer can drop the elements and link them to each other in order to create a workflow. The panel on the right (14) shows the properties of the currently selected element.

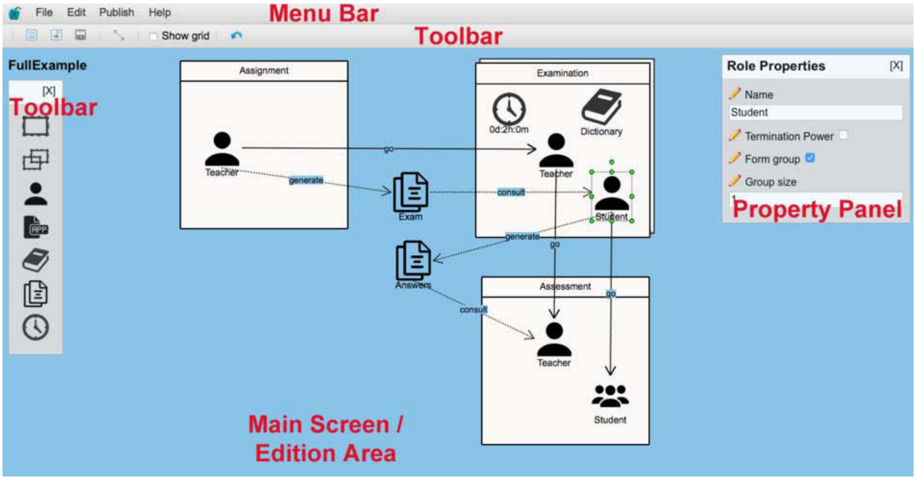

Figure 14. A screenshot of the simplified institution specification editor.

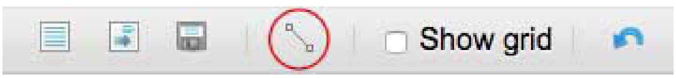

Figure 15. Toolbar that appears at the top of the main screen

Figure 16. Basic elements involved in the design of a workflow: activities, roles, resources, applications, files and timeouts

An activity can be created by dragging and dropping an icon representing an activity (first icon, from left to right, in Figure 16) from the toolbar to the central screen. In the same manner the designer can create group activities by dropping the group activity icon into the edition area (second icon, from left to right, in Figure 16). As mentioned in the previous sections, each role, application, resource or timeout in a workflow must be part of an activity or a group activity. Each of these elements can be added to an already existent activity by dragging and dropping the corresponding icon inside the activity or group activity. The icons that correspond to a role, an application, a resource and a timeout are the third, fourth, fifth and seventh icon in Figure 16, respectively.

The sixth icon shown in Figure 16 represents a file. A file can either be part of an activity (intra-files) or of the workflow (inter-files). Like the rest of elements before mentioned, the user can drop files into the main screen area, either inside or outside any of the already existing activities.

Activity order and roles in file sharing Arrows play an important roles in PeerFlow because they define the relationships between the basic elements of a workflow. The

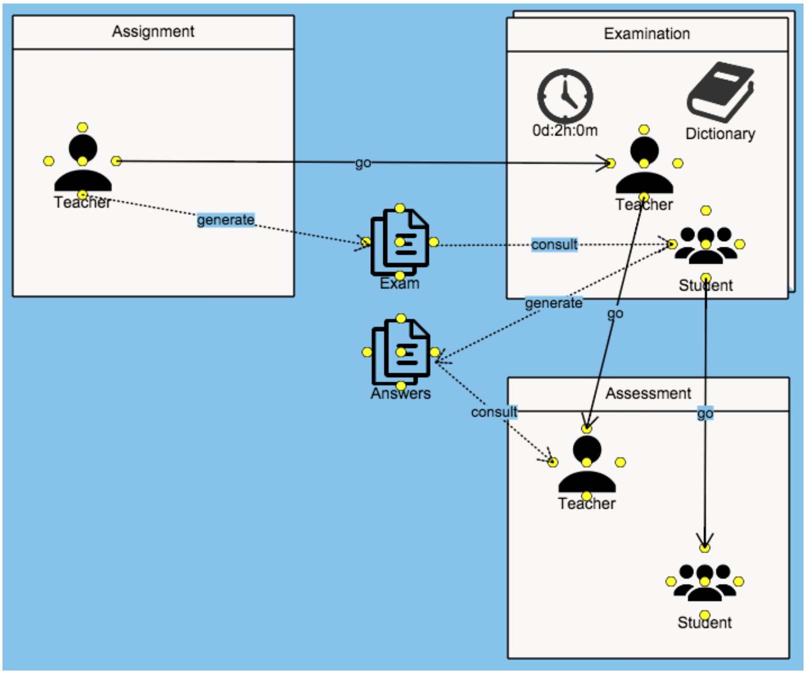

Figure 17. Glue points as guidance for drawing arrows between basic elements.

designer draws arrows between the roles in different activities to determine how the users playing those roles can move from one activity to another. Those movements determine the order in which activities are performed by each role. In a similar fashion, the designer needs to draw arrows between inter-files and the roles that produce or consume them. The PeerFlow Editor provides useful tips on how to draw arrows. The icon for connecting basic elements is the fourth icon in the toolbar that appears at the top of the main screen, which is highlighted in red in Figure 15. When the user clicks on this icon, each linkable element shows five yellow glue points, as illustrated in Figure 17.

All the designer then needs to do is to draw an arrow between glue points of the desired elements. In order to provide more expressiveness to workflows, the PeerFlow editor lets designers add labels to arrows. The designer can also modify default texts of labels by changing the label text input that is shown in the property panel when the arrow is selected. The editor establishes the default text of each label in line with the elements the arrow links and its direction. For instance, the label of an incoming arrow to an inter-file contains the text ‘generate’, which means that the role on the tail side of that arrow is in charge of generating the file.

3.3. The New Year’s Resolution Workflow

After we have explained how to create workflows using the PeerFlow Editor, we are ready to design a more ambitious workflow. In this subsection we present the New Year’s resolution workflow, which is aimed at improving students’ skills in using verb tenses when speaking about New Year’s Resolution. Figure 18 illustrates the New Year’s resolu-

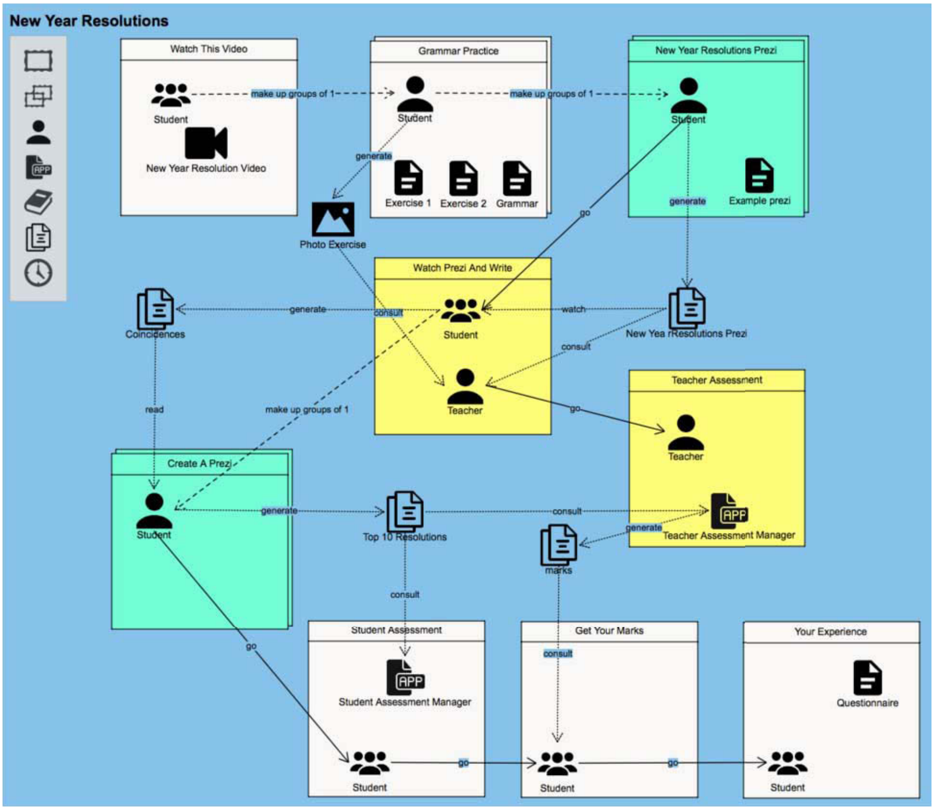

Figure 18. The New Year's Resolution workflow.

tion workflow. The new example consists of 9 educational activities, which are described next:

Watch this Video: This is the first activity in which students participate. In this activity, the designer requests students watch a video where people talk about their New Year's resolutions. We represent the video as a resource within the activity as can be seen in Figure 18.

Grammar Practice: This is the second activity in which students are involved. This activity should be performed individually by each student. The goal of this activity is for students to use the future tense. Students first should read the grammar text, and then do the exercises. When a student finishes, she needs to upload pictures of the finished exercises showing her marks. These exercises will be evaluated by the teacher only. The practice and the pictures will only improve students' marks, they are not mandatory. In the workflow the grammar text and the exercises are represented as resources. We also use the file resource named 'Photo Exercise' to capture the fact that students' results need to be exchanged between the activities 'Grammar Practice' and 'Watch Prezi and Write'.

New Year Resolutions Prezi: In this activity students learn how to create presentations using Prezi, a cloud-based presentation tool. As a result of this learning, each student creates a Prezi presentation about her New Year's resolutions. The presentations are accessible for the teacher and the other students in the next activity. In the workflow we include a resource called 'Example Prezi' in which students find the necessary materials to learn how to create presentations with Prezi. The file 'NewYearResolutionsPrezi' allows us to exchange students' presentations from this activity to the proceeding one.

Watch Prezi And Write: In this activity students watch their colleagues' Prezis and list the more frequent resolutions in a word file. We use the file 'Coincidences' to model the exchange of those word files between this activity and 'Create A Prezi' activity. This is the first activity for the teacher. The teacher receives the results the students obtained in the 'Grammar Practice' activity as well the New Year's presentations generated in 'New Year Resolutions Prezi' activity.

Create A Prezi: In this activity students individually create a new Prezi with the top 10 resolutions of the group. The new presentations are going to be evaluated in a collaborative way in the proceeding activities so we need to exchange those files among the activities. To do that we use the file 'Top 10 Resolutions' which is generated by every student and consulted by two applications in the 'Teacher Assessment' and 'Student Assessment' activities.

Teacher Assessment and **Student Assessment**: It is now time for assessing students. A Prezi generated by a student in 'Create a Prezi' activity is then evaluated by the teacher and other students in the 'Teacher Assessment' and 'Student Assessment' activities. The applications 'Teacher Assessment Manager' and 'Student Assessment Manager' manage the assessment process within these two activities. The 'Teacher Assessment Manager' sends students' presentations to the teacher. When receiving a student's presentation, the teacher can choose to assess the presentation. In that way, the teacher does not need to mark all students. The rest of the work is done by the 'Teacher Assessment Manager', which, as well as providing students' Prezi to the teacher, learns to assess students as the teacher does. On the other hand, the'Student Assessment manager' in'Student Assessment' activity asks students to assess the presentations of other students. In turn, a student may decide to assess, or not, a received presentation. The 'Teacher Assessment Manager' is also fed with the assessment that each student makes about other students' Prezi. Finally, the'Teacher Assessment Manager' proposes marks for students that are later validated by the teacher. The 'Teacher Assessment Manager' and the 'Student Assessment Manager' work together as a Collaborative Assessment tool as described in Section 5.

It should be noted the workflow designer also needs to define on which criteria the assessment of a presentation is based. Among the properties of the 'Teacher Assessment Manager' application, the designer specifies the rubric used to assess students as seen in Figure 19. This rubric file includes the assessment criteria for the workflow.

Get Your Marks: In this activity students are notified about their marks.

Your Experience: Finally, as a complementary task, students are requested to respond to a questionnaire aimed at collecting their experience using the PeerFlow framework.

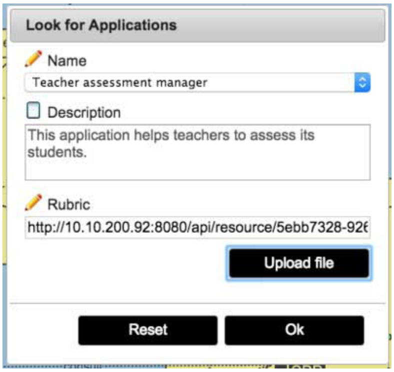

Figure 19. Defining assessment criteria as parameter of the Teacher Assessment Manager.

4. PeerFlow

In this section we explain how to activate a lesson and how students and teachers can participate in such lessons. For more information about the technical background we refer for example to [2,3,1].

4.1. Entering PeerFlow

To enter PeerFlow, you need to go to http://peerflow.iiia.csic.es and log in (or create an account if you don't have one).

Once logged in, you will see the screen displayed in Figure 20. This is the 'search' screen. On the top of this screen (as highlighted by Figure 21 you see a menu with five menu items:

- search
- running
- publish
- edit
- quit

The 'search' button takes you to the search screen. The search screen enables you to search for lesson plans that have been published, so that you can launch them (see Section 4.3).

The 'running' button takes you to a screen that looks almost identical to the 'search' screen. It also allows you to search, but this time the search will display lessons that are already currently running, rather than lesson plans that can be launched (see Section 4.4).

The 'publish' button allows you to publish a lesson plan that was created with the PeerFlow editor and that you saved on your hard drive. Alternatively, you may also publish a file directly from inside the PeerFlow editor (see Section 4.2).

The 'edit' button takes you to the PeerFlow editor where you can design a new lesson plan and save it to your hard drive or publish it, as described in Section 3.

Finally, the 'quit' button allows you to safely log out of PeerFlow.

4.2. Publishing a Lesson Plan

Let us now assume that you have already created a lesson plan with the PeerFlow editor and stored it on your hard drive (recall that to create the lesson plan, one needs to click the 'edit' button, as described in Section 3, and save the file).

When you have logged in to PeerFlow, click the 'publish' button. This will bring you to the screen displayed in Figure 22.

4.3. Launching a Lesson

Once you have published a new lesson plan you can launch it immediately. Alternatively, you can launch a lesson plan that has been designed and published by someone else.

In order to launch a new lesson, you first need to find the lesson plan among all lesson plans that have been published. In order to do so, click on the 'search' menu item. You will now see the screen displayed in Figure 20.

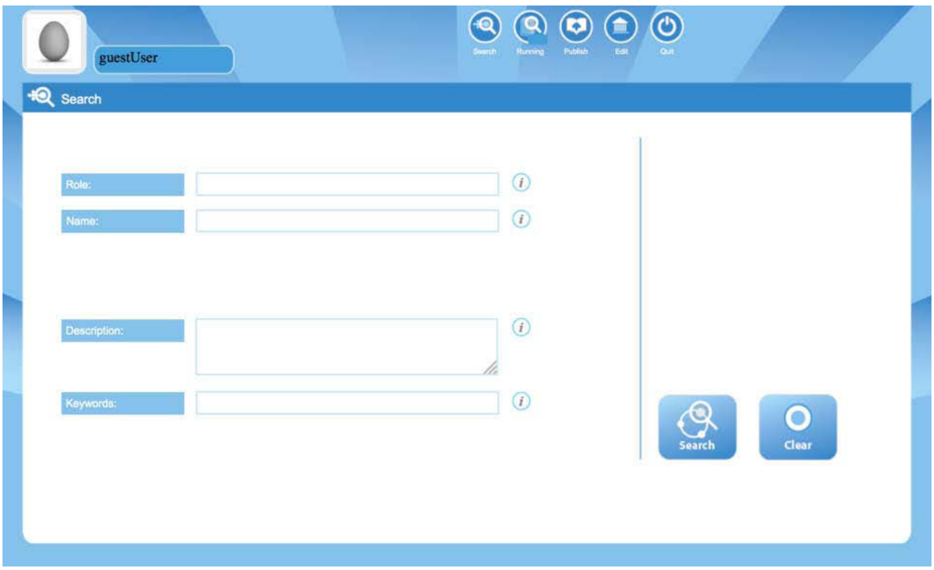

Figure 20. The search screen. This screen allows you to search for lesson plans that have been published.

Figure 21. The menu in the top of the PeerFlow screen, enlarged.

Figure 22. The Publish Screen. This allows you to publish a lesson plan that is stored on your hard disk.

In order to search through all published lesson plans you simply click the square ‘Search’ button in the bottom right of the screen. You may also apply a detailed search by filling out the fields provided by this screen.

For example, if you have given the name ‘EnglishCourse1’ to your lesson plan, then you can fill out that name in the ‘Name’ field such that the search will only return the lesson plan with that name.

Similarly, you can search for all lesson plans that fulfill a specific role, for example, you could limit your search to those lesson plans containing the role ‘student’. Similarly you can search for lesson plans with a certain description, or lesson plans that specify certain keywords. To clear the search fields again, click the ‘Clear’ button in the bottom right.

After clicking the square ‘Search’ button in the bottom right a new screen will appear that displays the search results; see the example displayed in Figure 23. The search results are listed in the left part of the screen. When you click on one of them, more details about that specific lesson plan will appear in the right part of the screen.

If the lesson plan you desire to launch is listed in the search results, you can select it, and launch it by clicking on the ‘Launch’ button on the bottom right. Another screen will then appear showing the name and description of the lesson plan you have selected to launch. You can now choose between ‘Launch’ and ‘Launch & Join’. The difference between the two is that in the case of ‘Launch’ you are just launching the lesson so that other users can participate in it, while if you click ‘Launch & Join’ you yourself will also enter the lesson. Of course, if you accidentally click ‘Launch’ while you wanted to Join it as well, you can still join it later (see Section 4.4). If you click ‘Launch & Join’ you will be taken to a screen (Figure 24) where you can choose the role that you wish to adopt in the lesson. Next, after clicking the ‘Join’ button a new window will open that will allow you to participate in the lesson, as explained in Section 4.5.

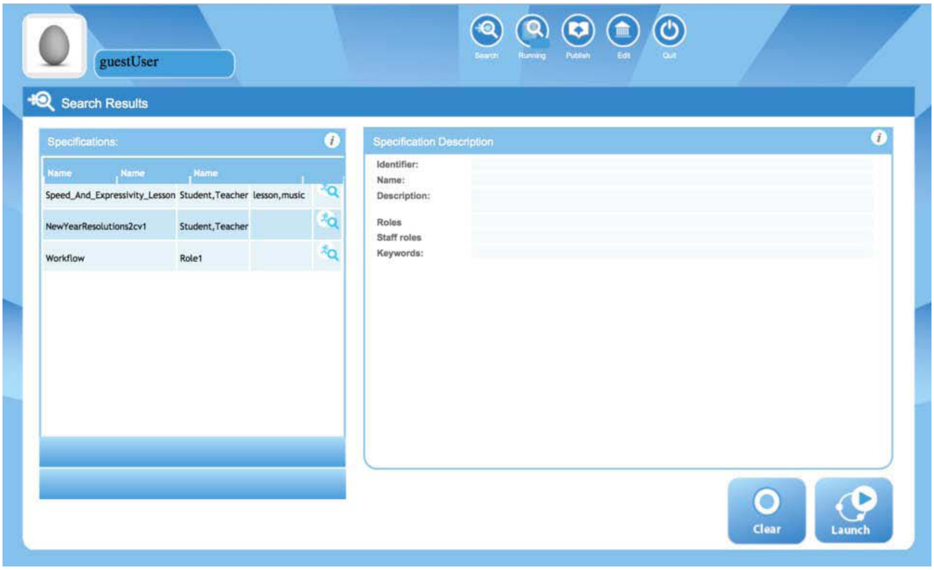

Figure 23. Search results. This screen displays the lesson plans that have been published and that match your search criteria.

4.4. Entering a Running Lesson

In this section we will assume that you want to participate in a lesson that has been started by someone else. For example, a teacher has started a lesson and you are a student and you want to take part in it. The teacher may have told you (e.g. by e-mail or in a physical class room) the name of the lesson, so that you can find it. In order to join the class you go to the PeerFlow website and click on the 'running' menu item. This will open a screen identical to the one displayed in Figure 20. You can apply exactly the same search criteria as when searching for lesson plans. When clicking the square 'Search' button in the bottom right, a screen such as the one in Figure 25 will appear. Note that it is almost identical to the one in Figure 23. The difference however, is that this time the search results will consist of running instances of lessons rather than lesson plans. For that reason, there is a 'Join' button instead of a 'Launch' button in the bottom right. After clicking the 'Join' button a screen appears where you can select the role you wish to adopt in the lesson, as displayed in Figure 24. After clicking the 'Join' button again, a new window will open that will allow you to participate in the lesson, as explained in Section 4.5.

4.5. Participation

Once you have entered a lesson (either after launching a new lesson and joining it as explained in Section 4.3, or after searching for a lesson that was already running as explained in Section 4.4), a new window with the participation screen, displayed in Figure 26, will appear.

This screen contains two main sections: the interaction screen on the left, and the action history on the right. Furthermore, there is a menu bar above the two main sections,

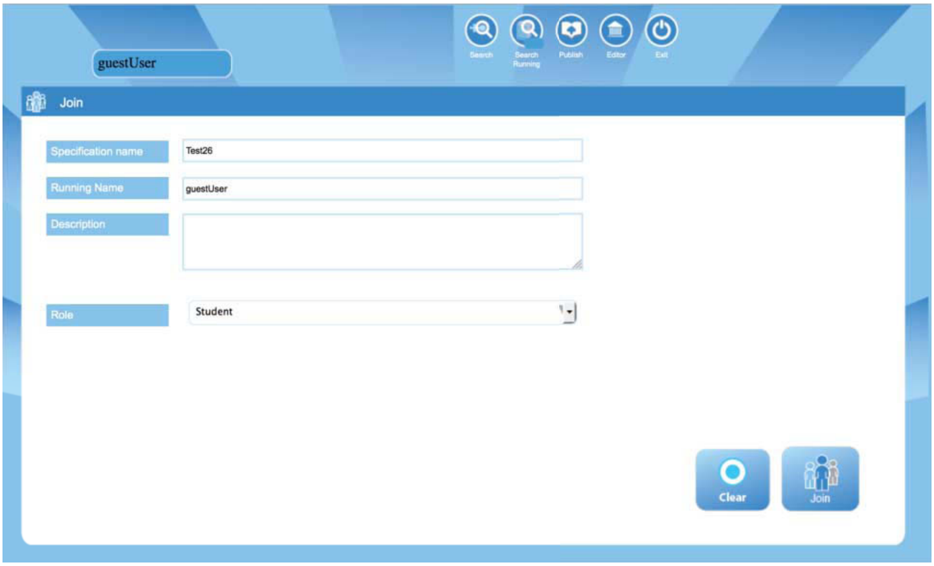

Figure 24. The Join screen. Here you can select the role you wish to adopt in the lesson (e.g. 'teacher' or 'student'. The roles you can choose between differ per lesson.

and in the top left you will find your personal avatar, your user name and the role you have adopted. The menu bar contains the name of the current activity in the left, and a list of menu items that you can use to move from one activity to the next. Some of these menu items may be disabled (or all of them). This happens when you are not yet allowed to move to those activities.

In the right of the menu there is a 'Map' button. When clicking it a map will appear that gives an overview of the lesson plan (see Figure 28), i.e. it displays all activities the lesson consists of, shows the paths you can follow to move from one activity to another and displays the actions that can be performed in the current activity.

The interaction screen displays a number of different icons: A circular icon for each participant in the current activity displays the avatar of the participants. As you can see in Figure 26, a number is displayed in the upper left side of the avatar. This number indicates the number of actions undertaken by this participant in the current activity. In the bottom right of each avatar two letters are displayed that represent the role of that participant.

Another circular icon represents a resource or a file (see Section 3 for an explanation of these concepts). In Figure 26 the icon is displayed as a book. Clicking on such an icon will open a new browser tab with the corresponding resource or file.

A cog-wheel icon represents an action you can perform. For example there could be an action named 'upload homework'. As a participant you can perform an action by clicking on its corresponding icon. When clicked, a pop-up screen will appear, as displayed in Figure 29, where the parameters of the action can be filled out.

Some actions may have *receivers*. This means that if you perform an action, some of the other participants in the activity will be notified that you did so, and will be able to see the parameters that you have filled out for the action. An example of such an action

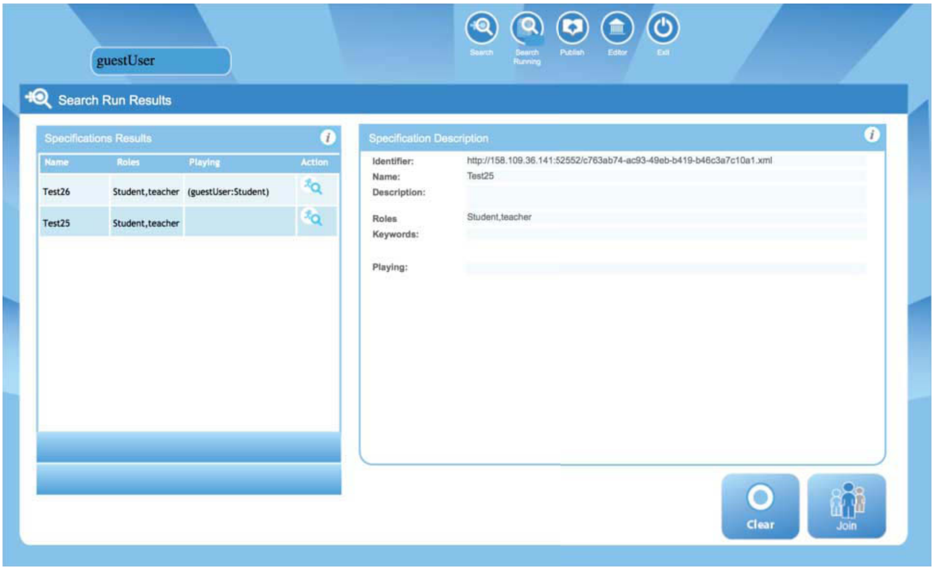

Figure 25. The results of a search for running lessons. This screen displays the lesson plans that are currently running and match your search criteria.

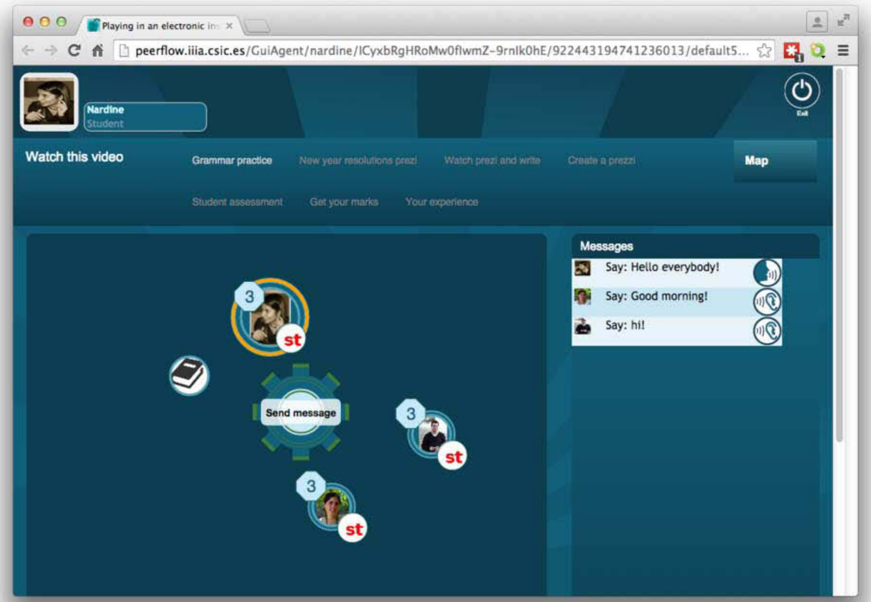

Figure 26. The participation screen.

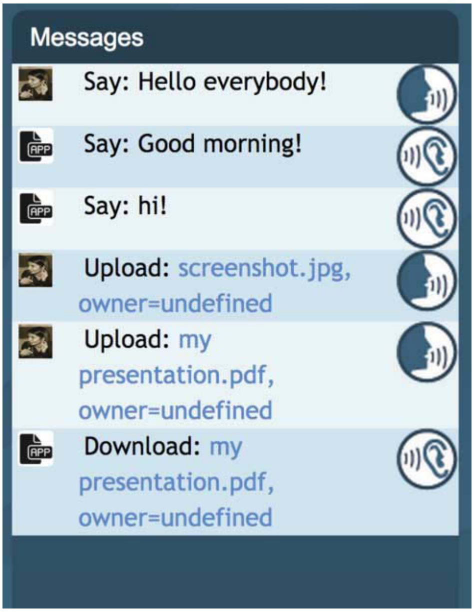

Figure 27. The messages appear in the right column.

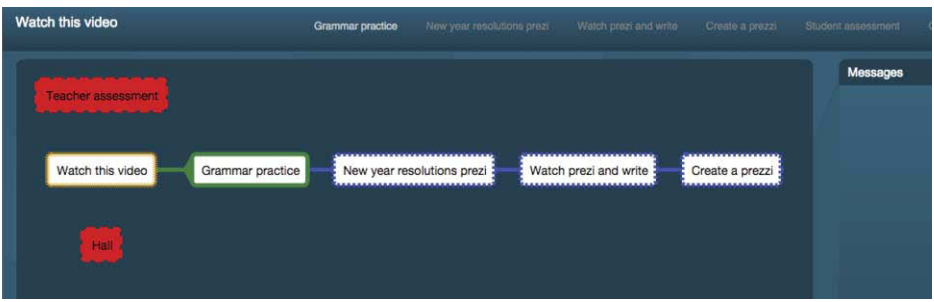

Figure 28. The map that shows the structure between the several activities of the lesson plan.

could be the action ‘assess’; when the teacher performs this action he or she would need to select a student in the activity as the receiver of the action and write the item from the assessment rubric that will assess that student’s outcome. The student will then receive a message with his or her assessment. Another type of action is simply a text message sent from one user into the activity of another user. The action history (Figure 27) displays a list of actions that have been either performed or received by you.

If you want to stop participating in the lesson there are two ways to realize this. One is by simply closing the browser window. This does not cause you to leave the lesson. You will still be in the lesson so the next time you log in to PeerFlow you can continue with the lesson. All you have to do is to search for running lessons to locate the specific lesson in which you were participating. Once you join, you will be back in the activity where you were when you left the lesson.

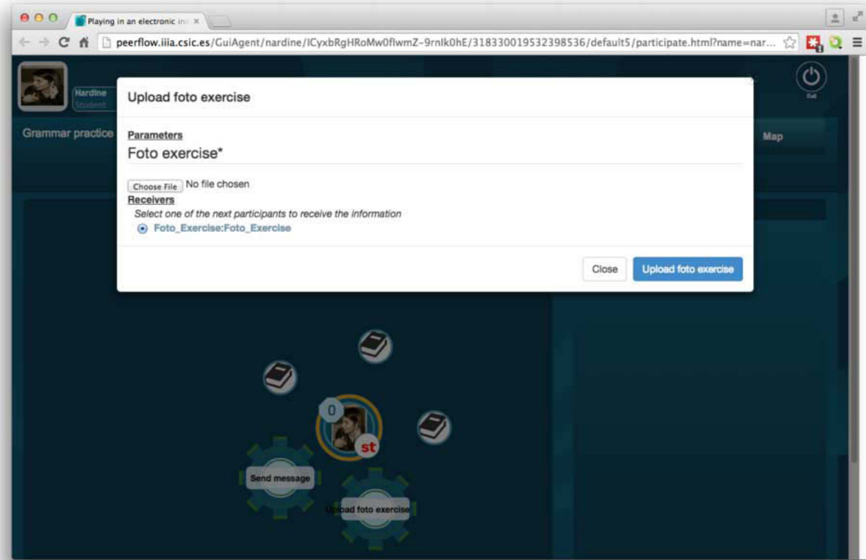

Figure 29. After clicking on an action a pop-up window appears to fill out the details of the action. In this example the student needs to select a file to upload.

The second way to leave a lesson is to click the 'Exit' button in the top right of the participation screen. Note however, that this will cause you to leave it definitively. If you change your mind afterwards and come back to the lesson this means that you will have to start all over again from the first activity. Moreover, it is not guaranteed that this is even possible, because some lesson plans may not allow students to join after the lesson has already started.

Although closing the browser window clearly has the advantage that you may come back, it is advised to use the other option if you really do not plan to come back. This is because if you do not exit the lesson but merely close the browser window, other participants in the lesson plan might expect you to still be with them, and hence they may wait for you for performing certain actions. By clicking the exit button you let everybody clearly know that you are no longer in the lesson and that others can continue the lesson without you.

5. Collaborative Assessment

In this section we provide details of the Collaborative Assessment (COMAS, short for Community Assessment) tool, which is integrated in PeerFlow and can be used as part of a lesson plan. This is an automated assessment service which (1) allows the tutor to evaluate students, (2) allows students to evaluate each other and (3) after analyzing students and tutor's assessments, generates automatic marks for students who have not been assessed by the tutor. The service is intended for intelligent online learning applications that encourage peer assessment, aiming to profit from these interactions to deduce new, meaningful information and reduce the work load for tutors.

5.1. Pedagogical Advantages of Peer Assessment

Self and peer assessment have clear pedagogical advantages [6,8,7,4,5]. Students increase their responsibility and autonomy, get a deeper understanding of the subject, become more active in the learning process, reflect on their role in group learning, and improve their judgement skills. Online learning communities encourage different types of peer-to-peer interactions along the learning process. These interactions permit students to get more feedback, to be more motivated to improve, and to compare their own work with other students' accomplishments. Tutors, on the other hand, benefit from these interactions as they get a clearer perception of the student engagement and learning process.

The method proposed here goes beyond current tutor-student online learning tools by making students participate in the learning process of the whole group, providing mutual assessment and making the overall learning process much more collaborative. Furthermore, peer assessment also has the positive, and much needed, effect of reducing the marking load for tutors. This is specially critical when tutors face the challenge of marking large quantities of students, as required by the increasingly popular Massive Open Online Courses (MOOC).

5.2. Method Description

Consider a lesson plan where students have submitted assignments that need to be assessed. Due to the large amount of assignments, the tutor is just unable to mark them all, although he or she can mark a subset of them. As part of the lesson plan, students are asked to mark each others' assignments and submit their feedback. The teacher will then be expected to mark a small set of assignments, and students will be expected to mark an even smaller number of their peers' assignments. Then, for all assignments not assessed by the teacher, COMAS will suggest marks based on aggregating students' marks. This aggregation takes into consideration the degree of trust the tutor has in each student's assessment (or mark). We define this trust measure based on the following two intuitions. Our first intuition states that if the tutor and the student have both assessed the same assignment, then the similarity of their marks can give a hint of how close the judgments of the student and the tutor are. Similarly, we can define the similarity of judgments of any two students by looking into the common assignments evaluated by both of them. We refer to these similarity measures as the *direct trust* between two people.

However, cases may arise where there are simply no assignments evaluated by both the tutor or selected students. In such a case, one may think of simply neglecting (or not taking into account) that student's mark, as the tutor would not know how much to trust that student's mark. Our second intuition, however, proposes an alternative approach for such cases, where we approximate that unknown trust between the tutor and the student by looking into the chains of trust between the tutor and the student through other students. For instance, we can say "if the tutor trusts student 1, and student 1 trusts student 2, then the tutor can likely trust student 2". We refer to this similarity measure as the *indirect trust*. And it is calculated using this approach in such a way that the longer the chain between the tutor and the student, the lower their indirect trust will be. To find such chains of peers, we build a trust graph from the history of assessments made, where nodes are the members of the community and edges join peers with direct or indirect trust relations between them.

The main task of the COMAS algorithm is to build a trust graph from the list of marks. Every time a new mark is added, the trust graph is updated with the new direct and indirect trust values. Once trust values are calculated/updated, final marks to be suggested are computed as follows. If the tutor marks an assignment, then the tutor's mark is considered as the final mark. Otherwise, a weighted average of the marks of students is calculated for this assignment, where the weight of each student is the trust of the tutor on that student.

5.3. *Example*

Figure 30 shows a simple example of a lesson plan with a tutor and two students where peer assessments are performed as part of the lesson plan. Assessments are made following 2 criteria: speed and expressivity, with an evaluation range from 1 to 10.

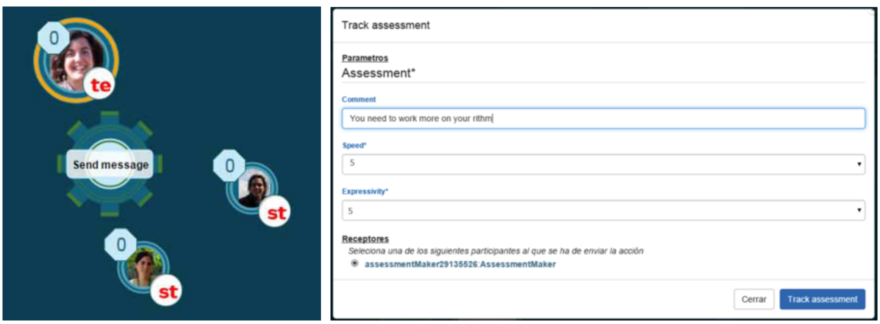

(a) Lesson plan in Peerflow (b) Assessment in Peerflow

Figure 30. Example of a lesson plan and peer assessment

Figure 31 shows the evolution of a trust graph from the history of assessments made. We write every assessment using the notation: $(exercise, assessor) = \langle speed\ mark, expressivity\ mark\rangle$. In Figure 31 (a), there is one node representing the tutor who has made the first assessment over the assignment ex_1, $(ex_1, tutor) = \langle 5,5\rangle$, and there are no links to other nodes as no one else has assessed anything yet. In Figure 31 (b), the student Dave assesses the same exercise as the tutor, $(ex_1, dave) = \langle 6,6\rangle$, and thus a link is created between them. The trust value between Dave and the tutor is high since their marks were similar. In Figure 31 (c), a new assessment by Dave is added for assignment ex_2, $(ex_2, dave) = \langle 2,2\rangle$. This has no consequences on the graph construction, since no one else has marked ex_2 yet. In Figure 31 (d), the student Patricia adds an assessment on ex_2, $(ex_2, patricia) = \langle 8,8\rangle$, that allows to build a direct trust between Dave and Patricia. The trust between them is low since their marks were not very similar. Also, an indirect trust between the tutor and Patricia is added, through Dave. This indirect trust is also low because, even though the tutor has a high trust on Dave, Dave does not has a high trust on Patricia.

The automated assessments generated when we reach stage (d) are:

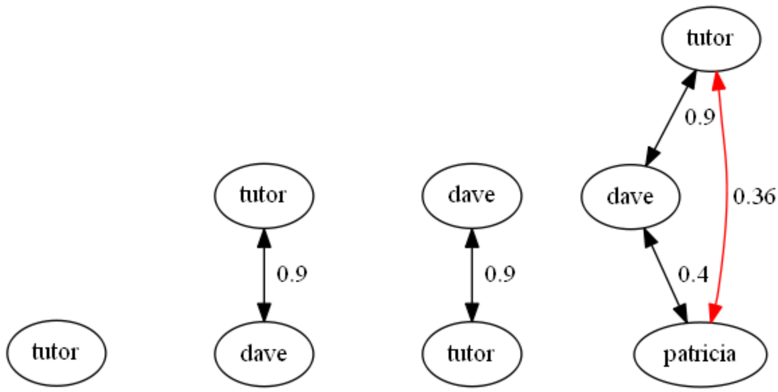

(a) $(ex_1, tutor) = \langle 5,5 \rangle$ (b) $(ex_1, tutor) = \langle 5,5 \rangle$ $(ex_1, dave) = \langle 6,6 \rangle$ (c) $(ex_1, tutor) = \langle 5,5 \rangle$, $(ex_1, dave) = \langle 6,6 \rangle$, $(ex_2, dave) = \langle 2,2 \rangle$ (d) $(ex_1, tutor) = \langle 5,5 \rangle$, $(ex_1, dave) = \langle 6,6 \rangle$, $(ex_2, dave) = \langle 2,2 \rangle$, $(ex_2, patricia) = \langle 8,8 \rangle$

Figure 31. Trust graph example.

- $\langle 5, 5 \rangle$ for exercise 1, which preserves the tutor's assessment
- $\langle 3.7, 3.7 \rangle$ for exercise 2, which uses a weighted aggregation of the students' assessments according to their trust relation with the tutor

6. Experiment results

6.1. Experimental setting

We evaluated the tools by running an experiment with real students. The experiment was carried out at a High School called Torras i Bages and located in L'Hospitalet de Llobregat, Barcelona, where Carme Roig works. Her students are familiar with the use of computers for learning, as she encourages them to use technology as much as possible. Students were informed that researchers from the IIIA research centre wanted to test a learning program. Out of a group of 25, 10 students were selected according to the following traits: familiarity with the use of computers, good behaviour in class, interest in the subject (English language), average knowledge of English (although this was not the most important point), all in all: "good students". The rest of the students were asked to wait for the next experiment. We were given a time slot of 2 hours to perform the experiment. The experience began with students logging in into the system and familiarising themselves with the PeerFlow environment. For a few minutes they were encouraged to click everywhere and to test what they could and could not do. None of them had problems with the log in process. Then they were asked to follow the lesson plan as it has been designed: they learned about the tasks they had to do, and they started performing them. That included watching a sample video of New Years resolutions, practicing their grammar by doing an online exercise, and doing their own presentation of their New Year's resolutions. At the same time they were able to talk to their friends online through the system, or they could ask the teacher online. At the end of the experiment they answered a survey, which we analyse next.

6.2. Survey results

The survey questions were simple and targeted to specific answers (mostly yes/no answers).

The first group of questions were related to the level of acceptance of the students with respect to the type of online activity performed *in the classroom*. Answers showed (Figure 32) a high level of acceptance towards the activity performed (Would you participate in this kind of online activities more times?: 100% acceptance, Would you like your classmates to participate?: 100% acceptance, Do you think evaluating you classmates can help you learning: 95% acceptance, Would you like to use always your PC in class and make the class activities with a computer online?: 95 % acceptance).

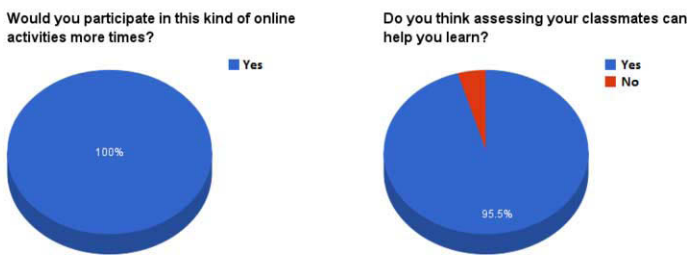

Figure 32. Level of acceptance of the type of online activity performed.

The second group of questions (Figure 33) were related to the level of acceptance of the students with respect to engaging in an online learning community where they could perform learning activities and about their habits of using online tools to study/learn. Students were more reluctant to engage in a learning community (Would you participate in an online learning community with other students: 63 % acceptance). Further questions showed that most students use online tools to communicate with each other for activities such as “make homework”, “ask doubts” or simply “talk about any subject”. When asked if they used any online tool to study/learn with their classmates a significant number (59 % of the classroom) answered yes and mentioned tools such as “edmodo”, “prezi”, “skype”, “google” or “traductor”.

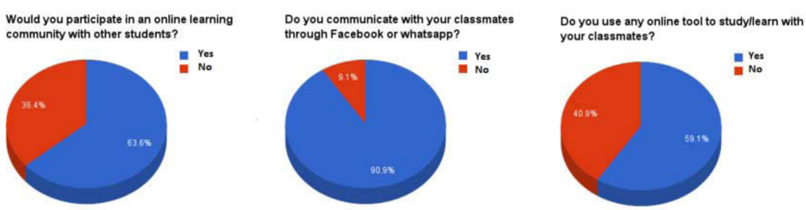

Figure 33. Level of acceptance in engaging in online learning communities and other online tools used to study/learn.

The third group of questions were targeted to the specific online application used in class, namely the PeerFlow tool. Answers showed (Figure 34) that most of the students liked the tool and were willing to use it again although many of them found it difficult to use (consider this is the first time they see and interact with the tool) and they considered the activity performed in class moderately difficult.

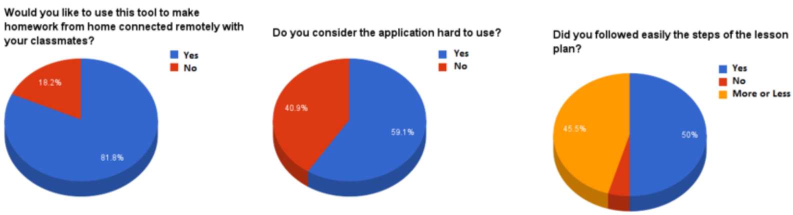

Figure 34. Level of acceptance of the PeerFlow tool. Difficulty in using the tool and difficulty of the activity performed in class.

In the last group of questions we asked students about their preferences about different features of the PeerFlow tool to improve its usability and to make it more friendly for potential users (questions included features such as the clarity/utility of the lesson map navigation tool, preference of navigation through the lesson plan map or through the menu, preference of choosing actions from the main window or from the lateral menu, preference of seeing the teacher's activities or not, login preferences, etc).

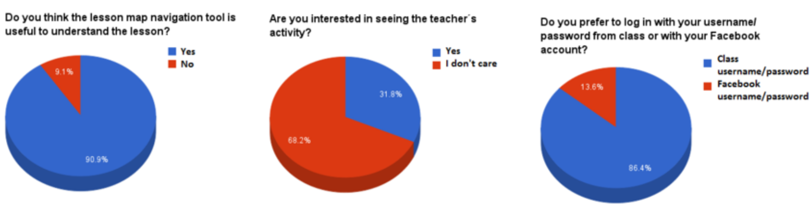

Figure 35. Student's preferences about different features of PeerFlow.

Overall, the survey showed an enthusiastic response from the classroom to the online learning activity performed and to the possibility of using again the Peerflow tool to perform more online activities with other classmates (more than 80-100 % acceptance). The response to the complexity of the tool and of the lesson plan performed was moderated, as it was the perspective of engaging in an online learning community with other students (about 50-60 % acceptance). Last but not least, we obtained a precious feedback from the students about specific usability questions that will help us improve our interface and functionality for potential final users.

7. Conclusions

This chapter has introduced the PeerLearn methodology and its associated tools. We design pedagogical workflows for students on the definition of rubrics (using PeerAssess) as the starting element that drives the creation of lesson plans (using LessonEditor). PeerAssess, an online database platform, helps teachers manage their groups of students, their lesson plans, and the scoring rubrics (or evaluation criteria). The web-based LessonEditor provides assistance to users to create workflows for their lesson plans. The six basic ingredients in creating workflows are: activities, roles, resources, applications, timeouts and files. Lesson plans created with the LessonEditor run over our web platform, PeerFlow.

We have also presented COMAS, the collaborative assessment tool, which is integrated in PeerFlow and can be used as part of a lesson plan. COMAS is an automated assessment service which allows each of the tutor and students to evaluate a small subset of assignemtns, and then, after analyzing students and tutor's assessments, generates automatic marks for students who have not been assessed by the tutor. This is specially critical when tutors face the challenge of marking large quantities of students, as required by the increasingly popular Massive Open Online Courses (MOOC).

Finally, we have evaluated the PeerLearn tools by running an experiment with real students. The experiment was carried out at a High School, Torras i Bages in L'Hospitalet de Llobregat, Barcelona, where 10 students from the English Language classroom were involved. By the end of the experiment, where students were interacting in a real lesson plan running over PeerFlow, the students had to answer a survey. The results of the survey illustrate an enthusiastic response towards using PeerFlow, and provide important feedback on usability that will help us improve our interface and functionality for potential final users.

References

[1] Dave de Jonge, Nardine Osman, Bruno Rosell i Gui, and Carles Sierra. D3.1 - electronic institutions for community building (v1). Public deliverable, The PRAISE Project (THEME 3: FP7-ICT-2011-8 no 318770), 2013.

[2] Dave de Jonge, Bruno Rosell, and Carles Sierra. Human interactions in electronic institutions. volume 8068, pages 75–89, Beijing, China, 01/08/2013 2013. Springer, Springer.

[3] Mark d'Inverno, Michael Luck, Pablo Noriega, Juan A. Rodríguez-Aguilar, and Carles Sierra. Communicating open systems. *Artif. Intell.*, 186:38–94, 2012.

[4] Valerie Hannon. *'Only connect!' : a new paradigm for learning innovation in the 21st century*. Centre for Strategic Education occasional paper ; no. 112, September 2009. Centre for Strategic Education, East Melbourne, Vic, 2009.

[5] Henry Jenkins. Confronting the challenges of participatory culture: Media education for the 21st century. 06 2009.

[6] Jingyan Lu and Zhidong Zhang. Understanding the effectiveness of online peer assessment: A path model. *Journal of Educational Computing Research*, 46(3):313–333, 2012.

[7] Karen Stepanyan, Richard Mather, Hamilton Jones, and Carlo Lusuardi. Student engagement with peer assessment: A review of pedagogical design and technologies. In Marc Spaniol, Qing Li, Ralf Klamma, and RynsonW.H. Lau, editors, *Advances in Web Based Learning , ICWL 2009*, volume 5686 of *Lecture Notes in Computer Science*, pages 367–375. Springer Berlin Heidelberg, 2009.

[8] Keith Topping. Peer assessment between students in colleges and universities. *Review of Educational Research*, 68(3):249–276, 1998.

[9] Grant P. Wiggins and Jay. McTighe. *Understanding by design*. Hawker Brownlow Education Moorabbin, Vic, 2nd expanded ed. edition, 2005.

Part IV

Experiences with Social MOOCs

This page intentionally left blank

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-177

Chapter 11. Using Social Media To Revive A Lost Apprenticeship Model In Jazz Education

Ed JONES [a], Harry BRENTON [b]

[a] *Leeds College of Music*
[b] *Goldsmiths, University of London*

Abstract.
The MusicCircle software was integrated into the BA Jazz course at Leeds College Of Music in a module titled Jazz Saxophone Performance in Context. In this chapter Ed Jones, a professional Jazz saxophonist and music educator, describes his experience using MusicCircle. Until recently, Jazz players were educated according to an apprenticeship model where young musicians ‘paid their dues’ by playing alongside older, established musicians. A huge amount of valuable knowledge was passed on in these environments. Jones used MusicCircle as a vehicle for reviving and rebooting this apprenticeship model by creating a blended learning environment where performance, theoretical and technical issues could be explored within a jam session / rehearsal environment. He found that MusicCircle increased student engagement and helped performers learn how to give and receive effective feedback.

Keywords. MOOCs, Jazz education, feedback support system, technology enhanced learning.

1. Introduction

MusicCircle is an educational platform designed for the creation of a community of learners [7]. We’ll showcase now how MusicCircle was integrated into a module entitled *Jazz Saxophone Performance in Context* within the curriculum of year 3 on the BA Jazz course at Leeds College Of Music. The aim of the module is to give students an opportunity to put into practice specific conceptual and theoretical approaches in a mixed seminar, workshop and rehearsal environment using a range of material and core repertoire. These concepts and theoretical approaches have an immediate practical application to the successful execution and performance of material associated with the genre. The module helps 3rd year students to prepare for their end of year recitals by allowing them to explore and develop a performance related approach that complements other theoretical and conceptual tuition.

The course was run by Ed Jones, a Jazz saxophonist/composer and educator. Ed wrote the following reflective account of his experiences using Music Circle.

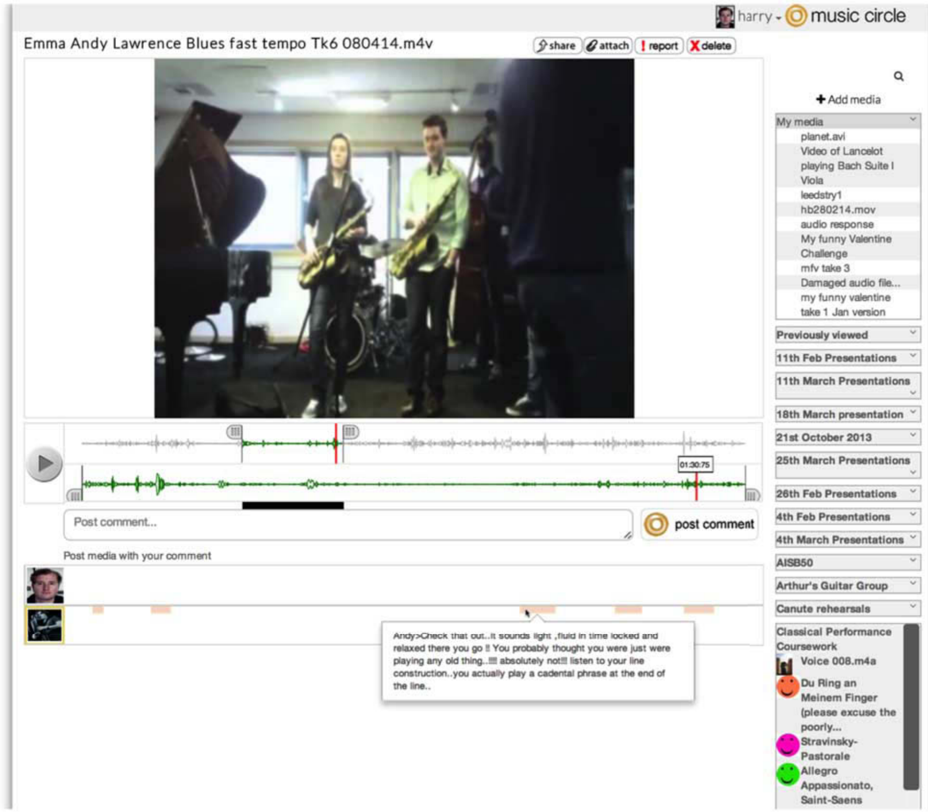

Figure 1. Jazz Saxophone: performance in context

2. Method

Until recently, Jazz players were educated according to an apprenticeship model where young musicians ‘paid their dues’ by playing alongside older, established musicians. A huge amount of valuable knowledge was passed on in these environments. This section describes an attempt to revive and reboot this apprenticeship model using MusicCircle.

Performance issues such as presentation, communication and successful program planning were also discussed and evaluated as part of an ongoing process of reflection and refinement. Another important aim of the module was to enable students to engage in critical evaluation of their process as well as developing abilities to engage and participate in peer feedback. The main participants were all 3rd-year undergraduate students with saxophone as their principal study listed in the examples as participants P1-10. To enable the process to be performance driven when needed, a further four participants were chosen from a pool of 3rd-year undergraduate Jazz rhythm section players.

MusicCircle was introduced at the beginning of the year in semester one by two members of the project: Dr. Harry Brenton and Dr. Mathew Yee-King, who travelled from Goldsmiths University to Leeds to demonstrate the capabilities of the site, train students, and facilitate the practical setting up of the community within MusicCircle. MusicCircle was used throughout the module as a means to document the process each week. Students would have an online document as to trace their progress through the module and the shifts in how this interacted with their artistic process. These sessions were documented either through video or audio recordings. I also wanted to both actively participate in the ongoing peer assessment and to facilitate the students' engagement with critical evaluation applied to their process.

Most of the weekly sessions were recorded and documented either on video or through audio files, then uploaded by myself onto the specific community site within MusicCircle. Over a 7 month period between October 2013 and April 2014 there were a total of 37 tracks uploaded and 425 comments.

3. Motivation for using MusicCircle

My aim in the module was to create a kind of environment where a multitude of performance, ensemble, conceptual, theoretical and technical issues could be exposed and explored within a working rehearsal or jam session situation. An equally important aim was to try to develop a community where there is not only a free flow of information and ideas, but where the members are encouraged and feel free to share their observations and experiences of each other's work. I strongly believe that one of the problem areas in Jazz education is that the dissemination of theoretical knowledge is very often given precedence over, or divorced from, the practical application of theory to practice (the *doing*). I used MusicCircle because I believed that it could strengthen the link between theory and practise. In particular, I hoped that it would help students to develop a more layered and in depth perspective upon their development by creating a lasting online record of the learning process within a community of peers and tutors. The constant challenge to 21st century Jazz educators is to find a balance between an African, aurally-based, experience-driven process; and a Western, theoretically-driven, education model. Hal Galper, a leading Jazz pianist and eminent and respected educator discusses this difference:

"The African teaching environment is the master/student relationship. [In this environment] the master coaches the student in both one to one as well as group contexts. The student learns the music [as well as] absorbs the teacher's attitudes [and their modes of thinking]. As mentioned in the Forward to 'Forward Motion', the student not only learns the 'whats' of music but the 'hows' as well. Seymour Fink, in his article 'Can You Teach Musicality' (May/June 1997 issue of Piano & Keyboard magazine), defines these two processes as 'conscious factual knowledge (knowing what to do)' and 'procedural knowledge (knowing how to do it)'. [...] The western environment of teaching is the classroom. Students are gathered together in one room where codified scalar, rhythmic and harmonic information are transmitted en mass. The teacher/student relationship is less personal than the master/student relationship. Contact with the teacher is minimal. In this environment the teacher very rarely gets to know the student's way of thinking, often to the detriment of the development of each student's individual voice and mind set. [...] Although masters of Jazz pedagogy, academic educators rarely possess the experiential knowledge that has been handed down from the Jazz masters through the apprenticeship system." (Hal Galper, *The Oral Tradition*, http://www.halgalper.com/articles/the-oral-tradition/)

Although I might take issue with Galper on his rather generalised assertions about academic Jazz educators, he highlights an important dichotomy in the difference between the traditions. Both approaches have their strengths and limitations. The focus of the African tradition can be a little one dimensional if the student only copies exactly what the master has shown him or her, and the western approach can, on the other hand, be too depersonalised by its focus on information dissemination. However, if we combine the positive aspects of both traditions, intellectual and experiential, then a more balanced and richer pedagogy can be implemented.

I believe that assured and successful Jazz education involves both the teaching of and practical interaction between:

1. The dissemination of conceptual/theoretical knowledge.
2. Development of core musicianship skills (technical, aural and written)
3. The practical application of the above 2 through practice, rehearsal and performance
4. Exposure to the key and significant artistic developments within the music within a time line lineage through recordings.
5. An understanding of the key social and historical contexts which influenced these artistic developments.

The module focused primarily on the first three of the above areas, although references were made from time to times on significant artistic developments and social and historical contexts. To achieve a more balanced pedagogy on the module, I deliberately moved away from the traditional handout written material, either in terms of notation or written explanations, to an aurally-based, discursive and demonstrative method with the emphasis on immediate practical applications. For example, giving a demonstration and discussion on voice leading (a line construction that focuses on outlining the harmonic movements of the form) followed by an example of how the concept worked aurally and how this could be applied as an approach to playing through chord changes in diatonic harmony-based material. This was delivered as an entirely aural and practical based exercise, one that the students could refer back to through out the year. Students were encouraged to abandon the method of learning material from written sources and instead encouraged to return to audio source recordings and learn material by a combination of ear and analysis.

4. The apprenticeship model

Up until recent times, a vital component of the Jazz education process was the ‘apprenticeship period’. This was when a young musician was given the opportunity to collaborate and gain vital experience from playing alongside older more established musicians through joining their working ensembles. A huge amount of valuable knowledge was passed on in these environments. One striking musical example is King Oliver’s advice to a young aspiring Louis Armstrong whilst the later was a young musician in his band: “*Learn how to play the melody before you do all of those fancy embellishments*” (Hal Galper, *Technique, part 2*, https://www.youtube.com/watch?v=NehOx1JsuT4).

In short, Oliver was advising Armstrong to pay attention to the melody line before he improvised or invented variations on it, as this would make his “*embellishments*” much stronger as they would relate to the core melody of the composition. This proved invaluable advice as Armstrong quickly developed into one of Jazz music's finest melodic improvisers. It also remains an important part of any aspiring Jazz musician's apprenticeship to be able to first improvise using the key melody notes before embellishing these in a multitude of ways involving a variety of different improvisatory strategies.

As 'sidemen' in the bands of older musicians, younger musicians were in turn given exposure to a larger community of musicians, larger audiences and the critical media through these environments. Eventually they gained enough knowledge, confidence and experience to form their own group. By this point they had often become known to public and critics alike. This apprenticeship was very much embedded in the career path of an professional Jazz musician. The term 'paying your dues' became synonymous with the apprenticeship period. The apprenticeship system of education was primarily an African-derived oral tradition in the master/student model, which can be illustrated by the experiences of the young John Coltrane studying the music of the older, more established, father figure of the pianist Thelonious Monk.

"He [Monk] would play anything, maybe just one of his tunes. He'd start playing it, and he'd look at me I guess, and so when he'd look at me I'd get my horn and start trying to find out what he's playing. He'd continue to play over and over again and I'd get this part and next time he'd go over it again I'd get another part. Then he'd stop to show me some parts that were difficult and if I had a lot of trouble he'd show me the music. He'd rather a guy learn without reading, you know, because that way you feel it better. You feel it quicker when you memorise it, when you learn it by heart by ear. And so when I almost had the tune down, then he would leave me to practise it alone and he'd go out somewhere; and I'd just stay there and run over the tune. Finally I had it pretty well and then we would play it together Sometimes we'd just get through one tune a day" (Lewis Porter [5]).

It is clear from the recordings of Coltrane (prior to his apprenticeship with Monk and those from during that time) that the developmental curve of his artistic vision was completely transformed by the experiences that he went through by learning and playing Monk's music. This kind of intense methodology and learning is difficult to replicate in a western educational environment, but elements of it can be illustrated and encouraged. Sometimes the transference of knowledge can be more oblique, though also vital in developing and nurturing inquisitiveness and establishing research methodologies. More recently the saxophonist Branford Marsalis has talked of these kinds of experiences in similar environments combined with an oversight on his musical experiences from a particular moment of his artistic development at the beginning of his career:

"You become a product of your environment. I was out there playing with my brother Wynton's (Marsalis) band: I'm an R&B saxophone player, I have no fuckin' idea how to play Jazz. Before that, I'm at Berklee and I'm listening to all of these guys playing all this fast stuff, and the question I had was, 'If all of this shit is so good, then how come it doesn't sound as good as the stuff from 30 years before?"

"And since I couldn't have that discussion with anyone, then I had to figure it out on my own, just asking questions, talking to Art Blakey, talking to Benny Golson, talking to Dizzy Gillespie. I would just say, 'what did y'all listen to when you were growing up?' And one of the constants was, 'church music and rhythm and blues.' But the modern Jazz guys of my generation basically did neither, with few exceptions. In Dizzy's day, you had 15-year-old kids playing in church bands or playing in rhythm-and-blues bands or swing-based dance bands, which were groove bands. But in modern times you got 15-

year-old kids learning 'Giant Steps' at Jazz camp. So I basically had to catch up to all of this stuff and learn. And it took a while."

"I was trying to play like Coltrane in Blakey's band, and one day Blakey walks by and says, 'what the fuck are you doing?' I said, 'I'm trying to play like Coltrane,' and he said, 'no, you're not!' and so I sarcastically said back to him, 'oh, so the best way to learn how to play like Coltrane is to not listen to Coltrane, right?'; and he says, 'well, let me ask you this: when Coltrane was your age, what the fuck do you think he was listening to, tapes of himself in the future? You dumb motherfucker!' And he walked off. And he left me with it. And that's the thing that was great about him. Whether it was because he was emotionally detached or because he just instinctively knew, who knows?"

"But the thing is, he understood—although he wouldn't have said it this way—that regardless of the kind of profession you're in, whether it's sports or music or whether you make typewriters, it don't matter. The two things that you have to develop on your own are cognition and intuition. Teachers have basically supplied the students with cognition. But in the manner in which they do it, intuition doesn't bloom. So when Art Blakey dropped that turd on my head, he walked away and left me to sort that shit out. So an hour later I walked up to him and I said, 'So when Coltrane was a kid, what was he listening to?' And he says, 'Ah, that's the question! Ask Benny Golson.' So I called Benny and Benny says, 'Oh, yes, young man, you might find this very interesting. Who do you think was Coltrane's first major influence?' And I said Charlie Parker. And he says, 'That wasn't it. It was Johnny Hodges.' I mean, who can put that together?"

"So then I had to hold my nose and start listening to Duke Ellington. Because, you know, I didn't come here for his shit. But if these old fuckers say that this is what it is, then alright, I'll have to endure this. So then, after about two or three weeks of listening to Duke records, suddenly you realise, 'Man, these cats are amazing!' Because it reminded me of what I learned how to do on R&B gigs. There were a couple of club owners in New Orleans who said, 'Man, you cats play too many fuckin' notes. And if y'all won't learn how to play this music the right way, we just won't hire you.' That gets it home. Note to self: less notes! And once you start getting into the idea of what R&B really is, then it's beautiful. But if your appreciation of music is always on the periphery of it, which means that your entire study of the music is totally based on harmonic analysis, then funk is a zero. You know, $F7^{th}$ for four minutes and 25 seconds is nothing, if your study of music is totally based on harmonic analysis. But if you can suspend that part of your brain and hear what they're doing and what makes it effective, then James Brown is suddenly the genius that he is." (Branford Marsalis, *Interview by Bill Milkowski*, Jazz Times 11/01/12)

Marsalis explicitly cites the experience of working with and learning from Blakey as a vital positive influence on his development as a Jazz musician. The example above also illustrates an important distinction between cognition and intuition as elements of educational and artistic processes. What the above dialogue illustrates is that Blakey understood that the younger musician needed to understand Coltrane's own influences in order for Marsalis to fully comprehend the process that Coltrane went through to eventually be able to arrive his own conclusions and thus be able to create an individual distinctive style. Blakey understood the process as undoubtably he had been through a similar process she he himself was a younger musician. As Hal Galper notes "*This is a music that teaches you through doing, the process teaches you the process*" (Hal Galper, *Musical vocabulary*, https://www.youtube.com/watch?v=b4kVUIpfTPU)

In the last decade or so these opportunities for such apprenticeships have diminished due to significant economic changes within the infrastructure of the business side of the music. There has been a greater shift to younger musicians forming their own projects and communities within the larger structure, that focus on the development and promotion of their own individual projects. Whilst this shift has many creative and economic positives for the young Jazz musician, a significant part of the knowledge base gained from being in these older environments and structures is being eroded because there are far few opportunities for younger musicians to gain working experience with older musicians on a long term basis. There is also a high turn over of 'next big thing' syndrome where a completely unknown new young Jazz musician releases his or her first album without being known to either audiences or critics, or even musicians within the community. There is currently a healthy discussion ongoing within the Jazz education community on how to address the erosion of the apprenticeship system.

It is clear that a vital part of a musician's ongoing artistic progress is the development of both an ongoing research methodology (and the practical application of this) and the development of critical faculties towards one's own artistic practice and that of elders, peers and others within the community. In the past, the apprenticeship model facilitated much of this work, through consistent live and recording work enabling these dialogues to take place. There was also much more opportunity for musicians to receive feedback from peers and elders through participation in recording sessions, live work, jam sessions, as well as developing the kind of critical processes necessary for personal artistic growth.

I have been fortunate to progress through these kind of apprenticeship environments and I am only too aware of this shift and its implications. Part of my own apprenticeship took place specifically with working in the bands of older more established musicians such as Clifford Jarvis, John Stevens and Dick Heckstall Smith. All of whom were older, respected, and established artists in their own right and able to pass on invaluable information and insights that were fundamental to my development. These experiences, as well as playing in bands of musicians from my generation and regularly attending jam sessions, helped me eventually to form my own artistic identity. Below I have cited a selection of the many gems passed to me (as well as some key pointers on which artists who to listen to) which really helped broaden my aural and historical contextual knowledge. Much of the advice still resonates with me today: '*Play less*' ,'*Play more*', '*Forget everything I said just play*' ,'*Stop showing off, technique isn't everything you know*', '*You really don't know enough ballads*', '*If you just play what you know you're not improvising*', '*Stop using the rhythm section as a play along record to showcase your latest meanderings*', '*Has anyone told you that you play consistently sharp, bloody well tune up before the gig starts*', '*What the fuck do you think you're wearing?*' And the best of all from Dick Heckstall Smith: '*Any musician who at some point doesn't think he or she is crap... is crap*'.

I participated in this project because I hoped that MusicCircle could help bring back elements of the apprenticeship model by enabling and documenting rich layers of social interaction and feedback. For the entire duration the 20-week course, students were solely reliant on the visual/audio documentation through the use of MusicCircle as there was a deliberate strategy to remove any written documentation from the process (which in my opinion can be a hinderance to aural assimilation). I firmly believed that Music Circle would not only be an effective way of documenting the information disseminated during the sessions and act as a reference space for this, but also a tool by which the students could develop perceptions of their own artistic process in a multitude of ways, thereby emulating some of the processes embedded within the apprenticeship model. Furthermore the documented sessions would be a permanent record that could be viewed and reviewed many times. The design of the interface (Fig. 1) meant that feedback was very specific because it was constrained to refer a particular moments within the performances and would thus be evidence-based. The fact that it would publicly available to everyone in the group would make it more democratic by allowing discussion and development to progress in a conversational style. By increasing learners engagement with their own work and that of their immediate peer group, my prediction was that this could increase a deeper and enriching aural/experience knowledge base. It was my belief that MusicCircle would fulfil these aims as its platform offers the opportunity for the participants to refer back to the experience and to be able to constantly revaluate their perceptions and those of others of the events, as well as developing critical and communication skills both in the moment and after the fact. A central objective of the project was to examine if MusicCircle could recreate positive elements of the apprenticeship model which I believed would help them to facilitate the development of critical analysis and feedback skills.

MusicCircle was used throughout the module apart from the very first few weeks as there were a number of practical and problems to solve such as setting up wifi in the practice rooms. Each session was recorded in its entirety using audio or audio and video, after which most of the sessions were edited into specific smaller documents relating to performances of material. These were titled ‘takes’ (for example *Out of Nowhere Participant 8 Tk1*) following recording session protocol. This was so to clearly define each performance as unique and to encourage the comparison of different takes by the same performer. As the discussion of conceptual elements were embedded into pre- and post-performance discussions and demonstrations, these were kept as part of the documentation to act as points of reference pertaining to the take and encourage a dialogue about the quality of feedback.

There were some student attendance issues throughout the module, as this was not a compulsory module. However, a clear pattern emerged quite early on of regular attendees and these were the students that ended up using the system and contributing the most to the process. Out of the group of ten initial participants this became a core group of five or six regular attendees. I was able to evaluate at key periods on how the system was being used as it was very clear who was commenting and when and how frequently. At the beginning and end of each session I gave clear indicators to the students that I was monitoring the process and giving feedback throughout on each audio/visual file and this information and their responses were vital to the process. Out of a total of 37 tracks uploaded, there was evidence of effective feedback in over 50% of the commentaries. 18 of the 37 tracks achieved 10 or comments; and whilst quantity is not always an indicator of quality, these were the tracks that contained the highest number of examples of effective feedback.

4.1. Analysis of MusicCircle data

The project had five objectives which are discussed below:

1. To identify examples of effective feedback

2. To examine how feedback is mediated by MusicCircle
3. To identify barriers to giving and receiving feedback
4. To propose ways to help students develop self assessment and feedback skills
5. To consider how well MusicCircle supports a revival of the apprenticeship model

4.1.1. Identifying examples of effective feedback

This objective had the following criteria:

- Evidence of appropriate communication from the person giving feedback.
- Demonstration of engagement and appropriate responses from the receiver.
- Demonstration of potential or actual learning outcomes.
- Evidence of dialogue, consisting of more than just a one-off response to a posting

Some of the most effective feedback came in the track *010414 Another you Participant 8.m4v* which accumulated 48 comments in total. In general the students are supportive and positive about Participant 8's performance and encouraging in a clear appraising manner. For example, this one from Participant 6:

"On the whole this solo is greaaat man. There are a couple of moments but on the whole it sounds so solid and controlled."

However after a while they become more a little more unguarded. Again from Participant 6:

"Love the note and rhythmic choices at the beginning of this solo but feel like the articulation is a bit jumpy in parts if you can hear what I mean?"

Here, Participant 6 is being supportive as well at the same time opening up a discussion on his articulation,and accurately highlighting an issue. Participant 8 replies:

"Yeah I definitely hear that. I think it was a combination of me temporarily being aware of over tonguing and trying to correct it and a tendency to give up at the end of a phrase. I had been attempting to censor my playing to form more coherent phrases rather than ramble. Thus the phrase ended up being quite clipped or just found nondescript endings. Have gone back to some of the phrasing things X (another lecturer) did with us, despite their rigidity I've found them a little useful just to break out of what I'm used to. Anyone have any interesting ways of thinking about phrasing?"

Participant 8 has clearly gone back to the highlighted area in the track and clearly heard and understood Participant 6's observation and reflected on this. After considering the feedback already received I reply to the question at the end:

"Think about language; the rhythm of words, sentences, syllables. Read prose or poetry, analyse the sentence construction, then think about how you can use this musically. Listen to your self speaking, try to match your natural verbal rhythm with what you play on your horn"

Although Participant 8 hasn't posted a reply, I have been in communication with the student and have confirmed that they found this whole post very useful in terms of confronting areas of weakness regarding articulation as well as giving them new ideas through the discussion posted. Elsewhere on the track there are some positive and effective comments on the progress made on tone. I begin the thread with a positive comment: "*I can already hear the difference in your tone. A lot more control in terms of the air column. The sound is richer and more complex... thats a big difference in such a short space of time.*" Participant 8 is quite candid and open about the tone problems here: "*Thanks! my tone really bothered me, I couldn't seem to find a point where it settled. I'd seem to* *build toward something I wanted then lose it again almost immediately. Don't know if any of you find a similar thing with tone practice but I find it a constant task to maintain any work that I have done, it doesn't seem to become natural for me.*"

To which Participant 4 replies: "*This is a similar problem I had when I found that my intonation and tone would just go a bit wild at the higher range of the horn. What helped me was some warm ups that included practicing overtones on bottom Bb and doing a exercise Bob Reynolds uses for long tones. Start on middle D then play the C sharp below then back to D then the C below etc. I would then repeat this but going up the octave into altissimo. Put the metronome at 40 and plenty of time, take it really slow between each note but don't overblow. It makes you really listen to your own sound and is kind of meditative. Also check out some George Garzone's masterclasses in Russia (or Poland) on youtube where he talks about tone and warming up, it's about 35 mins in. I don't know if any of this helps, but let me know if it does!*"

Participant 8 replies: "*Yeah that's great Participant 4 thanks. It certainly does feel meditative. Considering that, I found using the metronome a little distracting but practicing tone in that way, moving through the horn, I've found really useful. I found out some of the exercises at the beginning of the Rascher book as well and have gone about them in the same way.*"

This exchange is an excellent example of community dialogue at work. The students are being supportive whilst exchanging a multitude of ideas and references for further research. Participant 4 gives some effective feedback in quite a neutral non-judgmental style but gives plenty of helpful references and details of his own problem solving methodology. Participant 8 gives some effective feedback in return by commenting on his own experience of using the exercise Participant 4 has described and sharing another exercise that he found useful. The whole dialogue demonstrates that the students increasing awareness that all this work is ongoing and part of artistic practice and research is part of that on going process. Most importantly they are sharing ideas and experiences freely as peers. This is moving beyond traditional prescribed learning outcomes, in that the curricular outcomes are usually the result of objectives that are put in place to achieve them. Here we have evidence of dialogues and responses generated out of improvised events and improvised conversational flows; they contain learning outcomes, but ones not defined solely by either curricular or objectives.

4.1.2. Examining how feedback is mediated by MusicCircle

The majority of the tracks uploaded were recorded on video. This was strategically planned as I wanted the students to be able to visually reference their body language and verbal communication during performance to show how they came across to an audience from that perspective. There are specific performance related issues that video demonstrates very clearly. Also the process of watching (not just listening) their performances provoked some interesting comments that were elaborated on. For example in the track *010414 Another you Participant 8.m4v* I start by commenting: "*Good clear count in and click however you miss the pick up so the beginning sounds scrappy.*"

To which Participant 8 replies: "*It annoys me that I missed the top of the head both times!! Also I still find it a little hard to address an audience without a fear of either sounding sarcastic or pissed. As daft as it is, are there any ways of practicing this at home?*"

Here he acknowledges my comment with a recognition of my statement but uses the opportunity to open a discussion about what really bothers him about the excerpt, his presentation. To which I reply:

"First practice relaxation techniques such as slow diaphragmatic breathing. Some people rehearse what they are going to say on each gig. Think about the tunes you are playing and if you want to communicate anything verbally about them. This is more relevant when they are original compositions of course because then its personal. Remember an audience wants you to communicate to them... thats why they paid their money in the first place. So anything at all goes a long way. There are no hard and fast rules on what to say, and everyone finds their own way. Go to as many gigs as possible and note how effective the band leader is re his or her verbal communication skills. Pick an effective speaker, copy or use them as a role model."

Video is an excellent medium for highlighting several layers of information in regard to performance. In addition to documenting the audio of a live performance, video captures complex interactions of information which flow and fold in on each other to produce new and sometimes unexpected improvised dialogues and responses. These can be either out of responses to both visual or audio events and combinations of these in each moment. If this had just been an audio file I'm positive that Participant 8 would probably have been naturally concentrating much more on musical aspects. Here he is able to highlight and process a multitude of issues both musical and of visual presentation that are sometimes problematic to encompass in a single classroom session. MusicCircle accommodates these multiplicities and allows a variety of information strata to be disseminated, discussed and elaborated upon in ways that conventional classroom activity cannot possibly deal effectively within real time. The timeline allows both student and teacher to return to specific event areas and go into them at depth, thus in turn creating more dimensions to the learning experience. The text message format of social media forces feedback to be concise and direct and focused on specifics, but it also encourages dialogue between all contributors.

4.1.3. Identifying barriers to giving and receiving feedback

There are a number of barriers and issues to giving and receiving feedback. Criticism in our culture is something which whilst not a completely taboo subject, is rarely discussed in an abstract or constructive manner. There are numerous issues at play here. The student might feel a social pressure to be non confrontational and therefore holdback or moderate their opinions and observations out of a fear that they themselves will at some point be 'in the spotlight'.

"Despite the cognitive benefits of peer feedback activity, research identified both cognitive challenges and affective barriers of this activity on learners. Providing peer feedback is a cognitively demanding task for learners because they have to use their knowledge and skills to review, clarify, and evaluate other people's work. Especially, learners may not possess the domain knowledge or skills to provide useful and meaningful feedback as learners are often novices in the field. As such, they may provide feedback at a superficial level that does not lead to critical thinking of their own nor does it contribute to peers' learning. Affectively, students may have anxiety about giving feedback or little confidence in assessing their peers if they are not used to this activity, as they do not want to appear to be criticising peers' work. For example, Ellison and Wu (2008) found that college students were uncomfortable providing peer feedback on blogs. In ad-

dition, peer feedback may not be perceived as valid by the receivers as peer reviewers are usually not regarded as a 'knowledge authority' by feedback receivers, and, thus, learners refuse to take the feedback seriously. In addition, learners' peer feedback performance also varies depending on their characteristics, such as thinking style and level of academic achievement [6]. Lu and Law [4] found that learners' ability to identify problems in peer work and give suggestions was a significant predictor of the feedback providers' learning performance. Davies (2006) [1] found that 'better' students were more willing to criticise their peers than weaker students'.

However, as part of artistic practice, feedback from a variety of sources is an essential component in fostering growth and artistic development. Students, when face to face within a group situation, might have reservations about being totally candid in their responses to each other's work; and in turn might feel threatened or vulnerable when their weaknesses are revealed or discussed. In many respects, the students' use of MusicCircle reflects and exposes these very human issues in several particular examples from early on in the project. Giving and receiving feedback is clearly challenging, and one might reasonably expect that the process could possibly be more daunting by using MusicCircle—as there is a permanence to the documentation of that feedback. However, as the project progressed, it became clear that by using MusicCircle frequently the group of regular users adapted and became more comfortable with giving and receiving criticism. I believe this was because, through the use of the app, students became acclimatised to the social architecture of its environment. They became more comfortable with the fact that it wasn't a face-to-face real-time interaction, which allowed them to become more confident in expressing their own opinions and reflections.

On Participant 8's *Out of Nowhere Tk2*, Participant 4 comments: "*great tone man, it's got the essence of Getz and Zoot sims in there. Maybe some decoration in the head in the second A section and going over the last 4 bars in melody. Apart from that, sound great !*". Here, he is framing the constructive critic with two pieces of candy at the beginning and end. The important information in this is that the observation that the melody needed more decoration in the second A section and in the last 4 bars is clearly stated. The point is very valid. A melody that is repetitive in form needs small variations in order to give a sense of development over the form.

However, later in the project, the comments become more discursive, which reflects the group starting to relax with the issues of giving and taking criticism. The following dialogue illustrates this. Participant 4: "*When Participant 2 takes the head I am fairly unsure about what to play or if I should play. Participant 2 did some improvised embellishing when I played the head and wasn't sure if I had to do the same. This is something we didn't decide on before we had played and I was trying to be musically and contextually appropriate. I ended up not playing anything as I was over thinking.*"

To which I reply: "*Thats a valid decision and under the circumstances a good one.If in doubt don't play!! ,however some long guide tones would have worked and varying them rhythmically to fit Participant 2's phrasing*"

Later on on another posting on the the same track, Participant 2 posts: "*Here I am expecting Participant 4 to continue playing the melody which is why I felt I could improvise a little bit around it... but when he doesn't come in with the melody it just sounds really bad with us both improvising.*". This is a very direct discussion of a particular event that Participant 2 felt Participant 4 could have dealt with in as more constructive way by coming in with the melody.

Here is Participant 4's reply: "*I thought the opposite and figured Participant 2 was improvising a line to go into the melody as he did when we played the head before the solos. When I realised this wasn't happening, I started playing the melody and we both ended up coming in at the same time. All in all just a bit of a misunderstanding.*". Unsurprisingly, this is a different viewpoint of the same event. My comment in summing up: "*This highlights the need for saxophonists to grab more practice time TOGETHER!!! that way you can learn how to to build the musical dialogue and relationship.*"

By this point, both contributors have dispensed with the earlier tones of politeness and are having a frank open dialogue about what happened and how they can learn from the experience. Here it seems the usage of the platform has enabled them to articulate their thoughts on the events much clearer and more importantly in a much less reserved fashion. I believe these exchanges have really broadened their perspectives as they are now much more talking about the music as an object/subject by removing their own personal vulnerabilities through discourse.

They end in agreement, but after much healthy discussion. Participant 2: "*I think the head sounded better on the way in, when we weren't playing over each other too much... I don't think we had listened to the same recordings of this tune, judging on our interpretations of the head when we try to play in unison*". To which Participant 4 replies: "*I agree, despite both knowing the same standard, we have learnt / heard different versions and minuscule changes in the melody make all the difference to it sounding tight and in unison*".

At this point, I thought it was important to contextualize the discussion. I contribute this observation: "*The head was much better on the way in..here there is a lot of indecision.Participant 2 as a default you should think of yourself as 2nd voice when playing with Trumpet or Alto..like they are the main voice,and Participant 4 you need to lead like that. Only after this relationship has been established can you begin to to move to the middle ground. Check out a recent Pete King/Steve Grossman gig on YouTube..Pete bosses Steve (Who is a legend) and Steve knows he is 2nd voice and plays appropriately (Mostly!!)*".

4.1.4. Proposing ways to help students develop self assessment and feedback skills

How does self assessment support student learning?

- It enables students to take responsibility for their learning and positions the assessment task as being for learning.
- It allows students to engage in critical self-reflection on their own learning, that is an essential aspect of their academic and personal development.
- It provides explicit opportunities for students to reflect on their presentations, essays and problem sets, and allows them to critically engage with their work, identifying strengths as well as areas for further development.
- It enhances students' ability to self-regulate their learning by allowing them to assess their own work and understand how to close the gap from current to desired performance. This is crucial in developing students as independent critical thinkers.

"A key purpose of Assessment for Learning is to foster student development in taking responsibility for evaluating, judging and improving their own performance... These capabilities are at the heart of autonomous learning and of the graduate qualities valued

by employers and in professional practice.” (London School of Economics and Political Science, Teaching and Learning Centre, *Notes of Guidance*).

The process of self assessment is helped and encouraged by placing the student within a community and environment where it is encouraged, nurtured, and expected that they will engage with the process as a result of themselves engaging with the process of critical evaluation and peer feedback on the work of others within that community and environment. “*By peer feedback we mean a communication process through which learners enter into dialogues related to performance and standards*” [3].

How does peer feedback support student learning?

- It enables the development of critical reflection skills and the ability to give constructive feedback to peers
- It enables students to gain initial feedback on their work, and in a timely manner, that they can respond to in future assignments.
- It enables students to engage with assessment criteria and internalise them for application in their own work.

One of the big lessons for me as an educator has been learning when to mediate and when to stand back and let the student learn through the experience. This also applies to facilitating ways to help the student develop feedback skills. I have tried to chose my moments when to add to the dialogue, sometimes to help to clarify the issues being discussed and at other moments to give specific examples often with reference to particular recordings or to approaches and concepts to consider. MusicCircle has been very helpful in this way as there is a facility to add further documentation to the captured event in many different formats ranging from text to other audio and visual documents as well as links to other domains where information can be found.

I also tried to improvise responses as much as possible to give an unpredictability to the flow of the dialogue. As Jazz should be improvised so should meaningful dialogue and there are times where feedback needs to be unpredictable and spontaneous, it needs to ask questions as well as give answers. Although one might possibly expect that MusicCircle could in some ways weaken this unpredictability through its permanence of documentation, many examples showed that the feedback process would led to new and unexpected trajectories which in turn generated new strata’s of information. Throughout the process when there was evidence of effective feedback, I indicated that this was an example of how the process of feeding back should work. Part of my approach has been forged from my experience in my younger years when playing with older musicians where more often than not I would ask questions constantly. A lot of the time I received similar advice such as ‘*don’t think so much, just play*’ which is a valid viewpoint particularly if one finds oneself thinking too much whilst playing then the performer isn’t in the desired state of listening to their sound in the context of within the ensemble. It is important that the analysis and reflection and the dialogues that flow from the discussion of the performance need to occur after the event. If the performer is engaged with the analytical process when performing this can inhibit the flow of musical ideas.

Feedback given on MusicCircle is more specific than most feedback in the classroom because it is always substantiated with recorded evidence. MusicCircle also provides opportunities for reflection which allows sustained dialogues to occur much more frequently. Students can return to very specific event areas and really listen and consider the quality of the feedback they are giving or receiving. I actually believe that using MusicCircle develops feedback skills through intuitive learning: the doing or the experience of doing facilitates the learning process. The social media format forces students to be direct and concise in their communication, and if this is then translated back into a rehearsal situation (where ideally complex information needs to be disseminated quickly) the student learns how to give, receive and process information quickly. This has proved to be a very valuable contribution to their process. By regularly encouraging students to give and receive feedback this changes their experience from insular and inward looking to communicative and collective. It facilitates an understanding that Jazz music is both a subject and an object and, by its nature, collaborative and collective rather than personal. Feedback needs to be open ended not closed and not ending when the seminar or rehearsal finishes: it needs to be a dialogue and a continuous one. MusicCircle can help create the space to facilitate this, and by doing so students can experience the process of intuitive learning—the process teaches the process.

4.1.5. Considering how well MusicCircle supports a revival of the apprenticeship model

From the start of the project I was under no illusions that MusicCircle would be some kind of magical replicant of the key elements of the apprenticeship model. However it was clear to me that MusicCircle bore some resemblance to the environment of the earlier model by providing an environment that encouraged discursive dialogue. However, the difference was that these dialogues generated by the use of MusicCircle were much community driven rather than defined by purely the teacher/student relationship. Nevertheless, the more enlightened of environments that fostered my own development (even though often led by an elder) definitely encouraged collective artistic dialogue and discussion within them. By documenting whole events, I was able to dispense with the visual/reading elements that often block assimilation and practical understanding of conceptual and theoretical approaches. I was able to demonstrate alternative pathways and immediate applications there, safe in the knowledge that the students could return to the audiovisual documentation as constant reference points. This was similar to how older musicians used actual chord shapes to improvise from, as opposed to the more recent (and rather unhelpful in my opinion) chord scale model which can inhibit true improvisation. New dialogues then could unfold if necessary to expand or solidify the information. If we revisit the Monk-Coltrane example from earlier, it would be as if Coltrane had a constant reference as to what occurred during those rehearsals. An important part of learning the language of Jazz music is the study of pre-existing model references through transcription and assimilation into the students playing. This is an ongoing and more often than not lifelong part of artistic development. However, what begins with mimicry of particular sets of syntax can evolve into a personal assimilation and thus the beginnings and roots of a personalised style. MusicCircle enhances the assimilation process in many ways. It acts as a reference to a specific moment in the students development and allows a comparative process to take place; rather like the Jazz artists in the 1920's-1960's had as they were continually performing and recording, the documentation and feedback process for them was the process by which they learnt. The students that became regular participants and contributors clearly benefited the most, as I was able to increase the amount and quality of my feedback much more than I would have been able to achieve within the confines of a one hour session. This enabled them to access an environment of continual appraisal and assessment from the viewpoint of the elder or master musician.

5. Conclusions

In conclusion, the use of MusicCircle during the academic year of 2013-2014 has been incredibly enlightening for me as an educator on many different levels. Before I started using it, I was able to quickly grasp its obvious potentially as tool to facilitate learning in music. Until some of the students started to see the benefits of usage it was very problematic to give them reasons why they should contribute. However, by the end of the second semester, the use of MusiccCircle demonstrated marked improvements in the areas of:

1. *Performance presentation*: Earlier performances were notably lacking in confidence of delivery but showed a marked iin the latter stages of the project. I believe that the access to recordings on MusicCircle increased the students visual awareness of their presentation skills as well as being able to self reflect on the received feedback.
2. *Self evaluation and peer feedback*: After an initial period of stasis, core students demonstrated a marked improvement in giving and receiving effective feedback.
3. *Quality and effectiveness of performances*: There was also a marked improvement in these areas as students became acclimatised to the pressures of performing for their peers on a weekly basis. Towards the later stages there was a marked change in the level of engagement on their behalf with the rest of the ensemble. MusicCircle enabled the students to reference entire lessons and performances after the fact as well as mapping a personal timeline of development throughout the period of use. By both engaging with self evaluation and returning to recorded documents, the students began to engage with the process not just as an academic exercise but a process of artistic growth and development. Through its continual use over two semesters it became clear that core users and contributors had developed a strong community identity. One that was able to give and receive information very freely. Whilst the use of Music Circle was not the entire reason for this development, its role in creating the identity was vital.

MusicCircle has the ability to document, generate and combine complex relationships between strata of information which would not normally be possible just within the classroom environment. In this respect, the timeline is key to making this possible. This has a multitude of positive learning outcomes. MusicCircle itself has an inbuilt ability to be able create ‘folds’ (a variety of complex sets of information strata intersecting the event). As the influential philosopher Gilles Deleuze notes on the act of folding: “*It is not the line that is between two points, but the point that is at the intersection of several lines.*” [2]. MusicCircle, by creating these folds, initiates new events from the combinations of these intersections of information.

Furthermore, at several points during the project, I received a variety of feedback in the classroom and through MusicCircle itself that made clear that its use to document and reexamine technical, conceptual, or performance issues was determining the kind of direction the curriculum content should move towards. For example, it was very beneficial for the students to be able to document and reference different versions of the same material. That is, the process of using MusicCircle itself forced me to reappraise the structure of the curriculum I had devised for each semester. Personally, I found that it liberated me from overpreparing the sessions. Instead, in response to the events and dialogues documented on the MusicCircle, I developed a more flexible approach—which quickly became an important and invaluable tool for my own development as an educator.

All in all, MusicCircle creates a rich variety of new dimensions, giving birth to improvised dialogues and responses that would not be possible in the traditional classroom environment alone.

6. Acknowledgement

We would like to thank all the researchers staff and student who helped to develop the MusicCircle system as part of the PRAISE project (FP7-ICT-2011-8).

References

[1] Phil Davies. Peer assessment: Judging the quality of students' work by comments rather than marks. *Innovations in Education and Teaching International*, 43(1):69–82, 2006.

[2] Gilles Deleuze. *The fold: Leibniz and the Baroque*. U of Minnesota Press, 1993.

[3] Ngar-Fun Liu and David Carless. Peer feedback: the learning element of peer assessment. *Teaching in Higher education*, 11(3):279–290, 2006.

[4] Jingyan Lu and Nancy Law. Online peer assessment: effects of cognitive and affective feedback. *Instructional Science*, 40(2):257–275, 2012.

[5] Lewis Porter. *John Coltrane: His life and music*. University of Michigan Press, 1998.

[6] Marjo Van Zundert, Dominique Sluijsmans, and Jeroen Van Merriënboer. Effective peer assessment processes: Research findings and future directions. *Learning and Instruction*, 20(4):270–279, 2010.

[7] M. Yee-King, M. Krivenski, H. Benton, A. Grimalt-Reynes, and M. d'Inverno. Music circle: an educational social machine. In L. Steels, editor, *Learning about Music in the age of (social) MOOCs*, page Chap 7. IOS Press, Amsterdam, 2015.

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-195

Chapter 12. Improving Music Composition Through Peer Feedback: Experiment And Preliminary Results

Daniel Martín and Benjamin Frantz and François Pachet
Sony Computer Science Laboratory, Paris

Abstract. To which extent peer feedback can affect the quality of a music composition? How does musical experience influence the quality of a feedback during the song composition process? To answer these questions we designed and conducted an experiment in which participants compose short songs using an online lead sheet editor, are given the possibility to feedback on other participant's songs and can either accept or reject feedback on their compositions. This experiment aims at collecting quantitative data relating the intrinsic quality of songs (estimated by peer evaluation) with the nature of feedback. Preliminary results show that peer feedback can indeed improve both the quality of a song composition and the composer' satisfaction about it. Also, composers tend to prefer compositions from other musicians with similar musical experience level.

Keywords. music composition, experiment, peer-feedback,

1. Introduction

Peer feedback has become an ubiquitous feature of online education systems. Peer feedback consists in letting students or participants in a class revise, assess and more generally comment on the work of other students. This model is opposed to the traditional one in which students' works are evaluated only by a teacher. Peer feedback is acknowledged to bring many benefits [5] such as saving teachers' time as well as other pedagogical positive effects [7]. Nowadays, music students use on-line social platforms to learn and exchange music ([6] and [9]). At the same time, music education is changing as music schools are incorporating e-learning environments [1] . With the increase of online learning communities and MOOCS [11], peer feedback is becoming more and more popular.

Peer feedback is not only useful in pedagogical contexts, it can be also used in creative tasks. In music composition, collaborative composition has been addressed in several studies [2], including studies focusing on computer-based composition and collaboration through e-mail [10]. There are online creative communities in which music is composed collaboratively by several users [12]. There is an increase of the interest of musicians on participating in these music communities [8].

In those creative contexts, the following questions are legitimate: to which extent peer-feedback can affect the quality of a musical composition? What is the influence of the musical experience of the composers involved in this process? To address these questions we have designed a music composition experiment based on anonymous one-way feedback with no dialogue. Such a scenario differs from typical collaborative composition contexts in which composers work together hand by hand in a composition. The experiment is not aimed at being realistic or to propose a new tool for collaboration composition, but specifically to collect quantitative data regarding the relation between feedback, skills and song quality.

We focus on the role of peer feedback in music composition, specifically in *lead sheet* composition. A lead sheet is a representation of a simple song consisting of a melody and a corresponding chord grid. We propose an experiment in which peer feedback consists in suggestions of changes of certain parts of the lead sheet: specific notes or groups of notes or chords. These musical suggestions can be accompanied by a text explanation. Once a feedback is posted by a participant, it can be reviewed by the composer, who then decides to either accept it (and modify the lead sheet accordingly) or discard it.

Additionally to the sheer effect of feedbacks, we also examine the characteristics of the composer, commentator or judge of the participants. Indeed, having an extended experience in music composition might be seen as a prerequisite to write a nice song or to give useful suggestions. However, previous research showed that expertise might not be as critical as we could expect [3].

2. Description of the experiment

Participants are instructed to write a short composition using an on-line lead sheet editor [4]. Then they are asked to give feedback to another participant's composition, and finally they are asked to improve their own original composition using feedback posted on their composition. Participants are divided randomly in two groups: participants in the control group (G1) do not receive any feedback, and try to improve the song by themselves, whereas participants from the experimental group (G2) may use the feedback received to improve their own song. The existence of these two groups is ignored by the users so that the results are not biased.

As we are trying to assess the impact of feedback on the quality of a music composition, we need to estimate the *quality* of all compositions as well as their various variations during the experiment. To do so we use social consensus to determine the quality of a song: participants listen and are given the possibility to ”like” other participants' compositions. The quality of a song is then simply determined by the number of likes obtained for that song. In the next section we describe in detail each phase of the experiment:

2.1. Questionnaire

Participants start the experiment by answering 15 questions about to their experience in music, and more specifically in music composition. For example, they are asked how many years they have studied music theory, how many years they have been playing in a band, which style of music they like more, how often do they compose... etc.

2.2. Original composition

Participants then write a short composition using the online lead sheet editor. A lead sheet is a particular type of music score widely used in jazz, bossa-nova and song-writing, consisting on a monophonic melody and a chord grid. All compositions have a fixed length of 8 bars; participants are not able to add or delete bars, but they can choose the tempo and the time signature of the song. Participants fill the 8 bars with a melody and chord labels (e.g. Dmaj7, Em7...etc.). Figure 1 shows a screen-shot of the lead sheet editor.

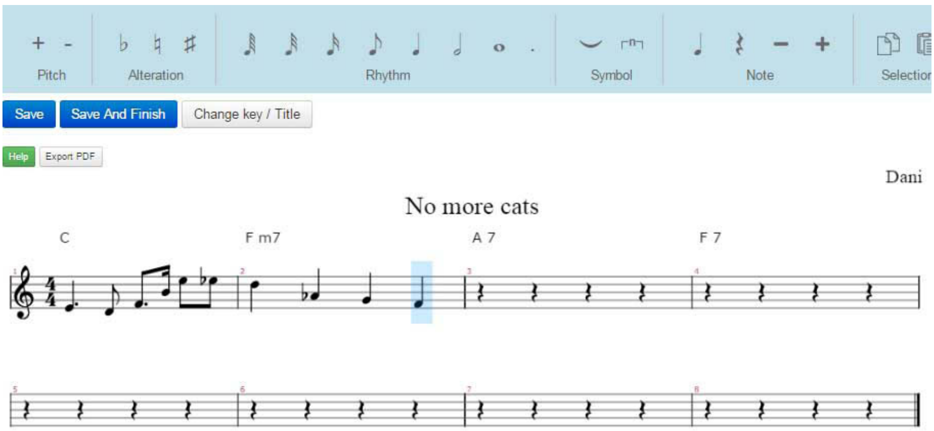

Figure 1. Screenshot of a composition being entered with the lead sheet editor.

Participants can listen to their composition with a basic MIDI player. When they are done they click on ”Save and Finish”. Next, they answer a questionnaire about their confidence in the quality, complexity and satisfaction on their composition.

2.3. Feedback Posting

Once they have finished their composition they are asked to give feedback to another participant by suggesting improvements in another participants' composition. Each suggestion can be at the most, two bars long. Participants can make as many suggestions as they want as long as they do not overlap. So, each participant can make a maximum of 8 suggestions (one per bar).

To make a suggestion, participants must choose the bar(s) to modify, then they can change the notes and the chord symbols. Optionally, they can also leave a text comment explaining their changes. Figure 2 shows a composition in which a participant is entering suggestions with an explanation. When they are finished, they answer a short questionnaire about their confidence on the suggestions they just made as well as their opinion on the original song they modified.

2.4. Improvement: Final composition

Next, participants are asked to reconsider their own composition and are asked to try to improve it. Participants from G1 (control group) are told that they unfortunately did

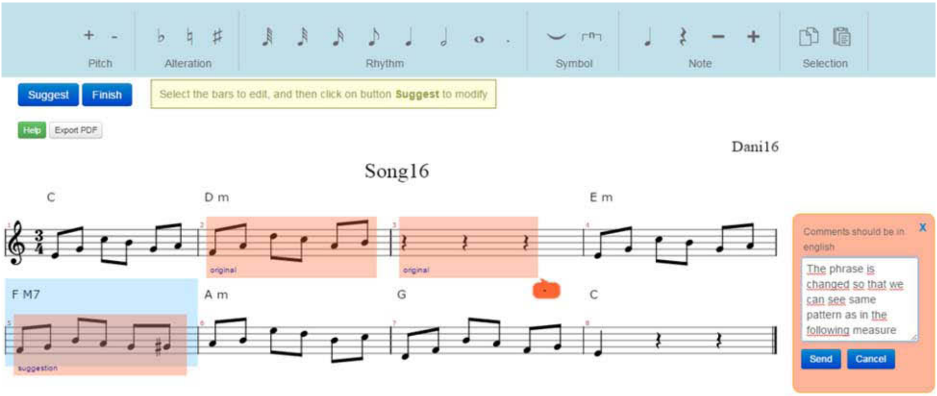

Figure 2. Screenshot showing a participant entering an explanation of the suggestion.

not receive suggestions and are encouraged to try to improve their own composition by themselves. Participants from G2 see the suggestions they received from two other participants. They can listen to all the suggestions. If they like a suggestion they can *accept* it, so that it is kept and the song is automatically updated accordingly. In addition to integrating suggestions, they can modify freely their composition. Once they are finished, they answer a questionnaire about their confidence on their own improvement and on their opinion on the suggestions received.

2.5. Evaluation phase

The last step of the experiment is to evaluate pairs of compositions from other participants. Each pair of songs consist on the original song and the improved song. Participants are asked to evaluate each song by place it in a vertical display with a legend from 0 ("I don't like it") to 100 ("I like it very much"). Participants do not know which is the original and the improved song when they are evaluating. One of the versions is presented as *song A* and the other as *song B* and this assignment is performed randomly. Participants have to evaluate at least 5 pairs of songs in order to finish the experiment.

3. Results

In this section we describe in detail the results obtained from each phase of the experiment.

3.1. Population

The experiment was conducted between February and July 2015. 66 participants completed the experiment (68% men and 32% women). Mean age was 29.2 years, ranging from 19 to 61. Musical experience was measured through a questionnaire with 7 items. The scale has a satisfactory sensibility with an observed range from 7 to 41 (out of 0 to

42) and we observed a mean of 28.7 with a Standard Deviation (SD) of 8.9. The intern consistency is satisfactory (Cronbach's alpha=.82).

Composition experience was measured through a questionnaire with 5 items. The results show an overall low level of experience concerning composition in our sample with a mean 6.9 (SD=6.1) on a scale ranging from 0 to 30). The intern consistency is satisfactory (Cronbach's alpha=.85).

3.2. Composition effects

Each participant was randomly assigned to either the control group (G1) or the experimental group (G2). No significant differences were observed between the two groups in relation to age, gender, musical experience or composition experience.

3.2.1. Composition evaluations

During the evaluation step, we checked if participants had listened to the songs before evaluating them. On the 1195 evaluations made, 219 were made without listening to the song. We removed those evaluations.

The songs were evaluated by an average of 8.8 different judges. The mean score of the evaluations made during the evaluation phase is 53.25 (SD = 13.26) on a scale ranging from 0 to 100. However, judges might be more or less strict, and some songs might have been evaluated by a particularly strict or generous participant. To take into account the severity of the judges, we have standardized the evaluations to get z-scores where the mean and standard deviation used are based on all the evaluations made by a given participant. As a result, the mean of the standard scores is approximately equal to zero, and a standard deviation of approximately .50. It should be noted that this final score correlates strongly with the raw score (r=.84). This result indicates that we had enough evaluations for each songs to avoid any severity bias.

3.2.2. Original Composition

The questionnaire that participants were asked to complete after finishing the original composition included self-estimation questions about the quality, complexity and satisfaction for their composition on scales ranging from very bad/simple/unsatisfied (0) to very good/complex/satisfied (6). We also asked them to evaluate the time they spent to make their composition and if they used an instrument to help them to compose (and which instrument if they did).

Results show a mean quality of 2.8 (SD=1.5), a mean complexity of 1.9 (SD=1.6) and a mean satisfaction of 3.2 (SD=1.6). Only the complexity is significantly different to the center of the scales which is 3 ($T(65)=-5.27$; $p<.0001$). This means that the participants tend to judge their work as rather simple (low complexity). We also observed positive and significant correlations between these three measures, ranging from r=.41 to r=.80.

During the suggestion step, we asked the participants to also rate the quality and complexity of the songs they had to comment. Each composition from the experimental group (G2) was commented by two different participants. In the end we obtained the score from the author and two other scores from two different commentators. Interestingly, there was no correlation between the scores from the original composer and the ones from the commentators ($r < .10$), but the two commentators did agree together on the quality (r=.80) and on the complexity (r=.70).

Moreover, from the judgments done during the evaluation phase (in which participants evaluate pairs of songs from other participants), the measurement of the quality of each original song (standardized to z-scores) allows us to estimate the composition skills level of its author. Surprisingly, we observed that the quality of the original song is only marginally related to the composition experience (r=.18, p=.15) or to the musical experience (r=.19, p=.12).

We also asked the participants whether they used an instrument to help them in their composition. Results show a marginally significant effect in favor of the use of an instrument on the mean quality score ($T(64)$=-0.87, p=.38).

The mean duration of the composition time of the song as evaluated by the participants is 30 minutes (SD=32 min) ranging from 1 minute to 240 minutes. This evaluation is largely underestimated by the participants because the real duration calculated from the time spent on the composition software is significantly longer (m=67 min; $T(65)$=4.20, $p < .001$). The correlation between these two durations is not very high, but significant (r=.46, $p < .001$) indicating that the error of duration estimation is not exactly the same for everyone. Interestingly, we observed that the quality of the original songs (from the evaluation phase) is not linked with the time spent to compose, whether it is subjective (r=.04) or objective (r=.03). This result suggests that in a situation where there is no time constraint, the amount of time devoted to compose has no effect on its quality.

Finally, there is a difference in the consensual quality of the original song, obtained from the evaluation of several participants (0.07 in G1 vs. -0.15 in G2). This could be due to differences in the group of judges evaluating each song.

3.2.3. Suggestions

In the questionnaire filled after making the suggestions, participants were asked how much do they think the song they are revising will be improved due to their modifications (on a 7 points Likert scale ranging from 0 ”very little”, to 6 ”very much”).

The participants from G2, the experimental group (N=30), received two suggestions for their final composition. Once they finished, we asked them if the suggestions received were interesting (on a 7 points Likert scale ranging from 0 ”very little”, to 6 ”very much”). Additionally, we recorded the number of suggestions they received and the number of texts comments received.

We ran a series of correlations between these measures and the improvement effect (the difference between the original song and the final song on the quality judgment score). None were significant, suggesting that neither the number of suggestions received nor the number of explanations for that suggestions have an impact on the improvement of a song.

3.2.4. Final composition

Overall, we can see that the control group, G1, does not improve significantly between the original song (m=.07) and the final song (m=.12) (improvement effect = .05, $T(35)$=0.94, p=.35). However, we do see a significant improvement for the experimental group, G2, between the original song (m=-.15) and the final song (m=.08) (improvement effect = .23, $T(29)$=2.47, p=.02). See Figure 3.

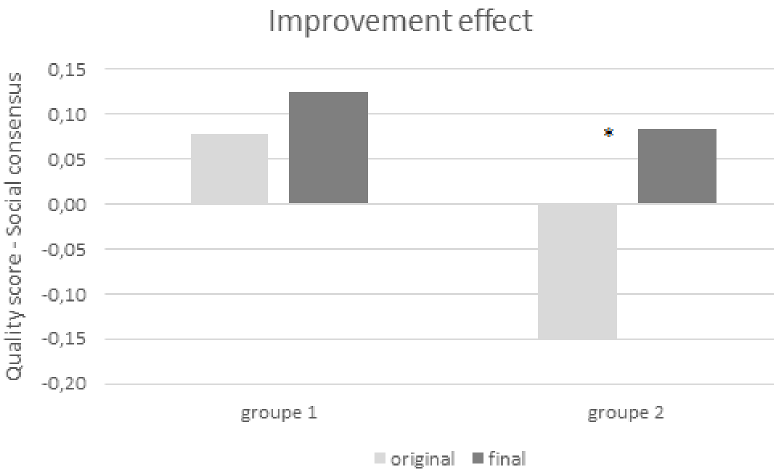

Figure 3. Difference between the original song and the final song on on the quality judgment score for the group without feedbacks (G1) and the group with feedbacks (G2).

We also examined the subjective evaluation of the participants concerning the improvement of their song. We constructed two composite scores. One from the self-evaluation scales of the original song (quality, complexity and satisfaction), one from the self-evaluation scales of the final song (quality, complexity and satisfaction). The intern consistency of those composite scores are satisfactory (the two Cronbach' alphas are above .81). We then conducted a mixed *between participants* (control and experimental groups) x *within participants* (original and final song) analysis of variance. We observed a significant interaction between groups and songs (F(1,64) = 7.07, p=.01). To explore this interaction, we used a post-hoc analysis with Tuckey HSD tests. Results show that participants who received suggestions had a significant improvement between the original and final song ($p<.001$) while the control group had no improvement (p=.49) See Figure 4.

When evaluating songs, users did not know which song was the original and which one was the final, as the order of the songs was determined randomly. This was a design decision to avoid the fact that participants could tend to rate better the final song, as it is supposed to be improved. Aditionally we wanted to ensure that songs were not better rated just because they had more modifications. To check this point, we used and melodic similarity algorithm [13] to estimate the similarity between each original and final songs. The correlation between the percent of similarity and the improvement effect based both on the composer's subjective opinion and on the scores from the judges are low (r=-.36, p=.003 and r=-.19, p=.13), which suggests that the improvement is not linked to the dissimilarity between the two versions.

3.2.5. Lead sheet editor

The software used was developed specifically for the experiment and we asked participant whether it was frustrating (0) or helpful (6) to compose with it. Results show a mean of 3.13 after the first composition and 3.41 after the final composition (the difference is

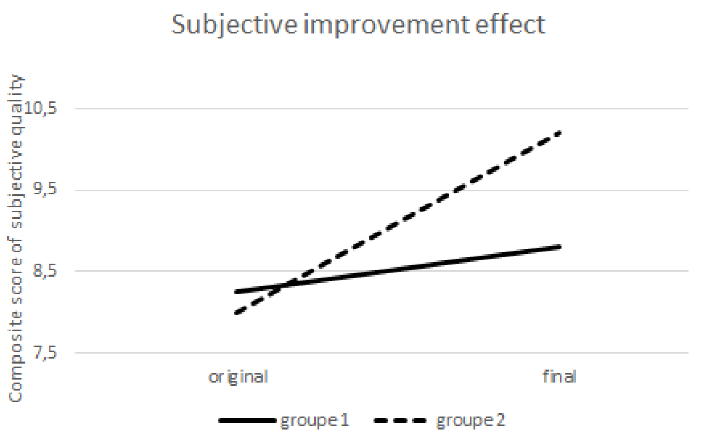

Figure 4. Self-esteemed quality of the original and final songs for the group without feedbacks (G1) and the group with feedbacks (G2).

not significant) which means that even if the online editor was not specially helpful, it did not hinder the composition process.

3.2.6. Experience effect on evaluations

To find out whether musical experience has an impact on the way participants judge song from other participants. We divided our sample of participants in two groups according to their experience as musician (based on the median). We also divided our sample of songs according to the experience as musician of their author. We then ran a two-way ANOVA to explore the effect of the experience of the judges according to the experience of the compositor. Results show a crossed interaction between these two variables ($F(1,61)=7.63$, $p=.007$) as illustrated in figure 5. These results indicate that experienced judges give high scores to songs from experienced authors and low scores to songs from non-experienced authors. It is exactly the opposite for the non-experienced judges. This means that participants tend to prefer compositions from other participants with similar experience. This could explain the difference in the evaluation of the original songs in G1 and G2. The groups of judges evaluating each song could have different level of expertise.

4. Conclusion

The aim of this experiment was primarily to examine quantitatively the impact of peer feedback in music composition and secondly to assess how important is the experience of the participants as musicians or composers in the whole process. Before any improvement or suggestions, participants had to compose an initial song. Interestingly, results show that participants' previous experience in composition did not impact the quality of their song. The same pattern was also found for the participants' previous experience as a musician. These two results suggest that the quality of a song (assessed here from social consensus) does not really tap in musicality but in something else, presumably

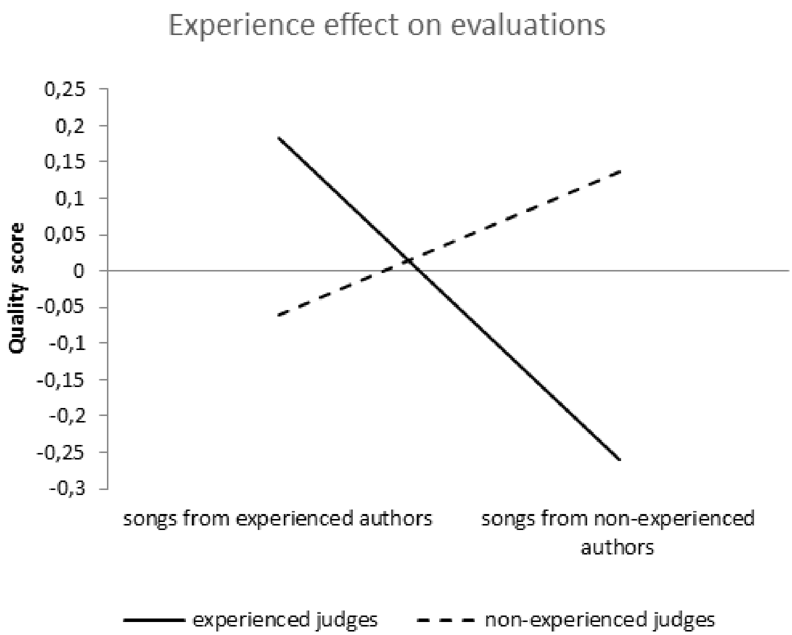

Figure 5. Interaction between the experience of the author and the experience of the judges on the quality score.

creativity. As suggested before, creativity might play a more important role than skills in this context [3].

Results show that composers who received feedback (G2) clearly evaluated better the improved song than the original, meaning that they were satisfied with the improvement they made. Further, the evaluation based on social consensus had a longer improvement also for G2. Hence, participants who received feedback not only felt that they had composed a better song after the improvement step, but they actually did. This basic finding suggests that improvements in music may be achieved even without real collaboration with dialogues and active interactions, but by simple suggestions on a single occasion.

Since there is a difference on the evaluation of the original songs between G1 and G2, we wanted to verify whether experience can make a difference when evaluating songs and we found out that participants tend to like more songs that are composed by other participants with similar musical experience.

Future directions of investigations include determining wth more details the influence of the participants' experience. This could be done, for example, by checking how song improvement relates to the experience of composers, commentators and judges. Further, we could assess more precisely which suggestions were actually used (or accepted) by the original composer to obtain a ranking of commentators whose suggestions are most accepted, as a measure of their pedagogical efficiency. We could check also if suggestions from experienced commentators are more likely to be used from inexperienced composers, or whether experienced composers usually accept suggestions of other composers, and how this impacts the improvement of the song.

References

[1] Pamela Burnard. Reframing creativity and technology: Promoting pedagogic change in music education. *Journal of Music, Technology & Education*, 1(1):37–55, 2007.

[2] Nicolas Donin. Domesticating gesture: the collaborative creative process of florence baschet's streicherkreis for 'augmented' string quartet (2006-2008). *Eric Clarke & Mark Doffman (eds.), Creativity, Improvisation and Collaboration: Perspectives on the Performance of Contemporary Music, New York: Oxford University Press*, forthcoming 2016.

[3] Michael Frese, Eric Teng, and Cees JD Wijnen. Helping to improve suggestion systems: Predictors of making suggestions in companies. *Journal of Organizational Behavior*, 20(7):1139–1155, 1999.

[4] Daniel Martín, Timotée Neullas, and François Pachet. Leadsheetjs: A javascript library for online lead sheet editing. In *First International Conference on Technologies for Music Notation and Representation (TENOR)*, Paris, France, 2015.

[5] Paul Rollinson. Using peer feedback in the esl writing class. *ELT journal*, 59(1):23–30, 2005.

[6] S Alex Ruthmann. Strategies for supporting music learning through online collaborative technologies. *Music education with digital technology*, pages 230–25, 2007.

[7] Philip M Sadler and Eddie Good. The impact of self-and peer-grading on student learning. *Educational assessment*, 11(1):1–31, 2006.

[8] Miikka Salavuo. Open and informal online communities as forums of collaborative musical activities and learning. *British Journal of Music Education*, 23(03):253–271, 2006.

[9] Miikka Salavuo. Social media as an opportunity for pedagogical change in music education. *Journal of Music, Technology & Education*, 1(2-3):121–136, 2008.

[10] Frederick A Seddon. Collaborative computer-mediated music composition in cyberspace. *British Journal of Music Education*, 23(03):273–283, 2006.

[11] On September. Behind the scenes with moocs: Berklee college of musics experience developing, running, and evaluating. *CONTINUING HIGHER EDUCATION REVIEW*, 77:137, 2013.

[12] Burr Settles and Steven Dow. Let's get together: the formation and success of online creative collaborations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2009–2018. ACM, 2013.

[13] Julián Urbano, Juan Lloréns, Jorge Morato, and Sonia Sánchez-Cuadrado. Melodic similarity through shape similarity. In *Exploring music contents*, pages 338–355. Springer, 2011.

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-205

Or should I say that MOOCS have been following me?

Yes, MOOCS are a hype. Yes, everybody is discussing if they are going to change the way we learn. Yes, I decided to follow one to see how it is and no, I didn't know what the acronym MOOC mean!

So the first step has been to look it up on Google ...

1) My Opinion Only Counts
2) Militia, Outlaws and Organized Crime
3) Massive Open Online Course

MOOCs are online courses on virtually EVERY subject; with millions of subscribers. I decided to follow the Songwriting MOOC of the Berkeley College.

It's Berklee! not Berkeley, BERKLEE!!!

Right; it's Berklee, a very famous College of Music in Boston, where many successful musicians have studied, such as ... well ... OK, the only reason why I know Berklee is that my boyfriend went there to study guitar when he was seventeen, instead of attending a very important exam in France.
The truth is that I don't know anything about music.

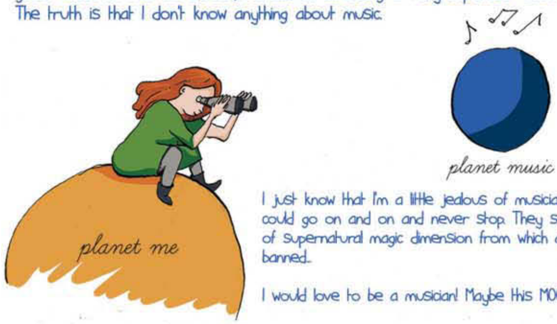

I just know that I'm a little jealous of musicians. When they play it looks like they could go on and on and never stop. They seem to be floating in some kind of supernatural magic dimension from which amusical humans are forever banned.

I would love to be a musician! Maybe this MOOC will make my dream come true!

So, I'm very enthusiastic about this MOOC and I think I'm going to learn very fast a lot of new things.
I signed up on COURSERA, the online platform for MOOCS, and I'm checking it everyday to see if the MOOC has begun.

So, I'm very enthusiastic about this MOOC and I think I'm going to learn very fast a lot of new things.
I signed up on COURSERA, the online platform for MOOCS, and I'm checking it everyday to see if the MOOC has begun.

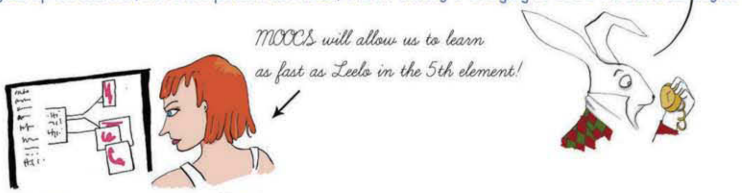

Finally the Songwriting course begins. The first lesson is composed by a few short videos - which are not presenting the course as one would expect. There is the Professor, a guy in his sixties called Pat Pattison. The first video is a very dramatic monologue about a … couple splitting up ?

He's very emphatic, just like and actor on stage, and he always wears a rock 'n' roll leather jacket.

He repeats things like 6 or 7 times.
I can't help but thinking that he's too "American". I mean, he's American like a bald eagle eating a hamburger at the White House.

But I'm impressed by the technical quality of the video. Everything is perfect!

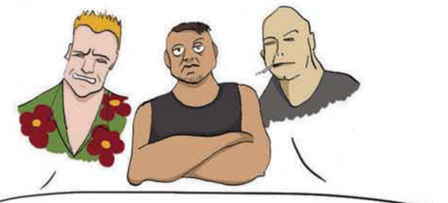

The course goes on and the theory behind it is beginning to unfold. The course is actually about how to write lyrics: Pat Pattison has some very clear rules:

Number of Lines	Length	Rhymes	Emotion
Even	Equal	Perfect	Stable
Uneven	Unequal	Not rhyming	Unstable

I'm of two minds about this. On the one hand it looks to me an over simplistic approach: did all great lyric writers really work in this way? I'm trying to picture Bob Dylan counting his verses to check if they are uneven to suggest a sense of instability and I just can't visualise it.

The whole thing seems too reductionist and too much based on "measurable" things, while hey, writing is not about measuring stuff!

On the other hand I begin to listen differently to lyrics, trying to analyse them.
Actually I begin to LISTEN, full stop.

For instance, on the bus!

The problem is all inside your head", she said to me
The answer is easy if you take it logically
I'd like to help you in your struggle to be free
There must be fifty ways to leave your lover
Fifty ways to leave your lover

long and articulated like a conversation...

Uneven number of lines and last line shorter : unstable! And it's true, you really wait for the song to move to the second half...

You just slip out the back, Jack
Make a new plan, Stan
You don't need to be coy, Roy
Just get yourself free
Hop on the bus, Gus
You don't need to discuss much
Just drop off the key, Lee
And get yourself free

This is stable... you even have two rhymes in the same line and lines are shorter and matching with each other ....

You actually feel like you HAVE to leave your lover and do it really fast!

(Paul Simon, 50 ways to leave your lover)

A MOOC is like a real course: not only you have to "attend classes" once a week but you also have classmates and homework.

I discover my classmates after one week or so: there is a huge online activity going on there! A forum with a many threads. People are coming mainly from the US and they are music lovers, many of them have bands and perform.
I read several discussions on the forums, but I don't feel like participating. Why? Maybe I'm the lurker kind, maybe Planet Music is still very distant from me...

Forum's web pages:

1 | 2 | 3 | ... | 45 | 46 | → Next

★ See top forum posters

No, I don't have a Mac.

As for homeworks, I have been doing some very easy quizzes every week, but the dreadful moment arrives in which I have to submit my work to PEER REVIEWING!

Now, not only making homework is in itself not very pleasant because it makes me feel like I was back at the elementary school, but I have a special problem with peer reviewing.

You don't know what you can expect from your classmates, you cannot evaluate their level of knowledge as you can do with a teacher, long story short: you cannot trust them! Yet you must submit your work for them to judge.

And by the way, hey, it's time to find an idea for our own lyrics!

We had to either choose from some proposed title or make up one of our own.
I've chosen the title "Two Story House" because I liked the fact that the word "story" can have two meanings.
Also, I have a lot of memories about a my grandparents' home where I used to went on holiday when I was a kid, and some kind of unresolved issues about it so... I thought that I could write the lyrics about it could be a way to solve them or at least archive them.

So, I went through four or five different version for the lyrics, here's the last one:

Two Story House

Grandmother called us but we were hiding
A summertime seldom seen
The pegs on the wire, the white sheets drying
Above the hill bold green

The house had two floors we had two stories
You dig " hide and seek , and still in the evening
you were playing out while I was dreaming
of sailors and warriors, dragons and fairies
Do you remember?

Then you kept reading and I went clubbing
just to seem a grown up
I lost myself in many bluffings
while you whirled up

The house had two floors we had two stories
You dig " hide and seek , and still in the evening
you were playing out while I was dreaming
of sailors and warriors, dragons and fairies
Do you remember?

We have found an agreement in not talking
to each other anymore
but sometimes I look back dreaming and hoping
That you remember

The house had two floors we had two stories
You dig " hide and seek , and still in the evening
you were playing out while I was dreaming
of sailors and warriors, dragons and fairies

I do remember

I was scared the first time I had to read peer comments about this intimate stuff!

But I was relieved when I read reviews like these:

REV #1: This is great. Could be a duet. Good imagery and good scene setting. I am surrounded by snow currently, but the opening l ines take me there. Nice work.

REV #2: Love the detail and the set up in the verse.
He first line in the chorus is great. Love the idea and how it sets us up for the rest of the song. My only problem (?) is with lines 2 and 4 of the chorus. I would have liked a rhyme of some sort not necessarily a perfect one or even a family one but some kind. It almost threw me off only cause it was specifically supposed to be stable.
Of course, lol what do I know? Overall I loved. Good work!!

REV #3: Very good. Nice unstable and stable structures here, and I like your use of a spotlight in the last line of the verse. I really like how descriptive the verse is -- you're setting the scene quite clearly here. The idea in the verse strikes me as stable, so I'd be inclined to not use it in an unstable structure.

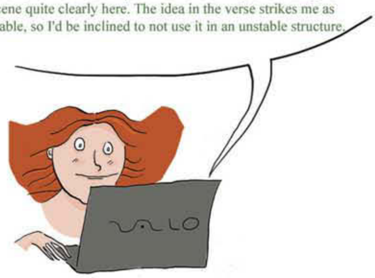

At the beginning this made me feel like I had to return the favour to my fellow students trying to tell them something really meaningful and helpful about their work, as far as I could not just it's good or it's bad.

My reviews of others' work are not saved on the Coursera platform, so I cannot copy them here. Anyway it was in the same style of my reviewers #2 and #3.

At the same time, I'm more and more interested in what is going on in the Forums, and I see that I'm not the only one who has a problem with peer reviewing, as since the course began the forum's threads are filling up with bitterness :

Typical threads at the beginning of the MOOC:

Hello from sunny Arkansas! (10 posts)

Hi from Chattanooga, can we be friends? (16 posts)

Yay! Looking for collaboration (12 posts)

Threads when peer reviewing has begun:

I am not satisfied with the reviews (53 posts)

I know who you are, Reviewer #2 (421 posts)

Post here if you hate your reviewers (683 posts)

Around the half – two thirds of the course, many things have changed:

1) I realise that a MOOC actually takes a lot of time and I'm struggling to find the time to attend classes.
2) Even if now I (kind of) trust peers reviewing me, I'm always uncomfortable when I have to read their comments
3) I'm not so enthusiastic about peer reviewing myself. I forget to do it at least four times
4) Quizzes and classes are not simple any longer. I have to repeat them several times before getting it right
5) I changed my mind about Pat Pattison

Once again I'm not the only one in the course. Berklee College has been carrying out a survey about what students liked and disliked about the course:

Dislike

Like

My admiration for Pat Pattison reaches a climax when I see him in action in a real class at Berklee College. He's coaching a musician who is playing and singing her own song. Without changing her lyrics, with just with some advices about how to stress syllables and other details he's making her song different. Much better

The last courses about rhyme types and prosody are dense and they get more and more difficult to follow.

At the end of the course, I'd crown Pattison as poet laureate :

But I'm not able to do the last assignments. The problem is that in the last two courses something new has come in my way: ACTUAL MUSIC!

I mean, I should compose a melody to match with my lyrics. Pat Pattison gives clear directions for that too, but after fantasizing a while about writing a brilliant song I find an excuse and skip the assignment.

I'm still thinking about doing the course twice to make it to the end.

Will I do it? I'm not sure, I think I'll need a practical (and unlikely) reason to do that, like being hired as a lyric writer by someone from distant Planet Music.
I also would like to back to the lyrics I wrote and change many things...
I feel like now, a few weeks later, I would look at it with a clear mind.
In the meanwhile, hey, I've got my Statement of Accomplishment to share on Facebook!

and it's signed !

This page intentionally left blank

Part V

Looking Backward and Looking Forward

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-213

Chapter 14. A Half Century Perspective On The Role Of Computers In Learning And Teaching

Ken KAHN
University of Oxford

Abstract. For more than fifty years people have been exploring how computers might enhance learning and teaching. The malleable nature of computers has enabled suggestions that a computer can act like flash cards, personal tutors, textbooks, reference books, virtual laboratories, quizzes, virtual spaces, lecture halls, and study groups. Perhaps the most radical suggestion has been to see the computer as something learners can creatively mold into something personally meaningful that is dynamic, interactive, and shared. And that the process of constructing such computational artefacts is rich in learning opportunities. These range from a deeper understanding of the subject matter of the constructions to high-level skills in thinking and problem solving.

Keywords. Constructionism, technology-enhanced learning, Seymour Papert, history of computers and education

Introduction

The idea that computers can play an important role in learning and teaching is over fifty years old. This chapter describes the history of attempts to use computers to support and enhance learning from a personal perspective. Instead of a complete history it attempts to highlight groundbreaking and significant ideas and computer systems that have led to today's efforts to provide technology-enhanced learning. Some systems use the computer to emulate older paper-based technologies. Others attempted to give the computer the role of teacher or tutor. The systems that are described in the most detail in this chapter are those that attempt to use the computer to provide novel learning experiences that were impossible or impractical before.

The 1950s through the 1970s were dominated by "computer-aided instruction" systems that attempted to teach in a very didactic and mechanical manner. These were based upon behaviorist theories of learning. Research laboratories at MIT, Xerox PARC, and the University of Edinburgh were exploring a very different approach. Instead of the computer programming the student, the student was given tools for programming the computer. Creativity and exploration were emphasized. Early attempts at computer tutoring systems were made. Programming languages designed specifically for learners were developed.

The 1980s saw the wide-spread dissemination of personal computers and programming languages for children. There were efforts to enhance these languages with new ideas from computer science. Media creation was combined with program creation. Intelligent tutoring systems were demonstrated to work well for a limited number of topics.

The 1990s saw the wide-spread use of "multi-media" to enhance education. Programming languages for children expanded into new territories. Learners were supported in building computer games and programming robots.

The 2000s saw the integration of the web into educational software. Learners were connected by the World Wide Web and able to easily share their constructions. Multi-user three-dimensional virtual spaces became popular places to explore their potential to enhance learning. Many explored the benefits of each learner having their own personal computer.

The 2010s saw the introduction of MOOCs (massive open online courses), web-based programming environments, and mobile devices.

1. 1950s and 60s

Alan Perlis saw the potential of computer programming for learning by science, mathematics, and engineering students in the mid-1950s. He began teaching the first freshman course on computer programming in 1958. In 1961 in a lecture at MIT he said "The purpose of a course in programming is to teach people how to construct and analyze processes [1]. J.C.R. Licklider commented "... I see computer programming as a way into the structure of ideas and into the understanding of intellectual processes that is just a new thing in this world".

In 1964 Kemeny and Kurtz introduced the Basic programming language, the first programming language designed for learners and beginners [2]. It contained many comprises due to the hardware limitations of the day. Variable names, for example, were limited to one letter followed by digits. While initially limited to use in universities, Basic became very popular with schools and hobbyists in the 1970s and 80s.

In 1967 Seymour Papert, Wally Feurzeig, Cynthia Solomon, and Danny Bobrow developed the Logo programming language. Unlike Basic, which was designed to provide the minimal language that can support student programming, Logo was designed to be a rich and powerful language. Logo is the result of "child-engineering" the best ideas in computer science at the time. It borrowed very heavily from the Lisp programming language which was being used by artificial intelligence researchers. Logo was conceived of as both a tool for learners to use to express themselves creatively and an "object to think with" [3]. Initially the projects created using Logo focused upon word and list processing and mathematics. For example, children constructed programs that generated poetry. By 1969 Logo was enhanced to control "floor turtles", robots that could be commanded to move forward or turn. This became the basis of the very successful turtle graphics when "screen turtles" were introduced in 1972.

A very different trend that began in 1960 is "computer-aided instruction". This was pioneered by the Plato system [4]. The Plato system initially focused on presenting multiple-choice or numeric questions and automated responses. Its initial innovations were in computer graphics and display terminals that it pioneered. The Plato system grew over time to include interactive simulations, educational games, and discussion forums.

But unlike the efforts around Basic and Logo, Plato was based upon a didactic teaching method instead of the programming languages' support of learner-centered problem solving and creativity.

Figure 1. The PLATO System

The idea of using computers in education was very radical in a period where computers were few and very expensive. As Hal Abelson, one of the earlier pioneers of Logo programming, said "You really have to try hard to get into the mindset of that time, because a computer in those days was something that cost several million dollars. And the idea that you would take the most advanced computing research equipment around anywhere, and you would let fifth graders ... start playing with it, it was just mind boggling. For the first 10 years of that, people just thought we were nuts" [5].

2. 1970s

The next decade saw substantial progress in efforts around the Logo programming language. Since the center of this research was the MIT Artificial Intelligence Laboratory it is perhaps not surprising that many efforts attempted to connect Logo and AI. Gerry Sussman [6] and Ira Goldstein [7] produced systems that helped debug and teach Logo. Danny Hillis wrote about AI projects that children could do in Logo. Radia Perlman developed special hardware to provide interfaces appropriate for very young children to construct Logo-like programs [8].

Figure 2. Radia Perlman's Button Box for Preschoolers

This was the decade when the concept of object-oriented programming was incorporated into programming languages for children. Smalltalk 72 and 76 were designed for children and inspired by Logo. (Smalltalk 80, however, was developed as a tool for professional programmers.) Director was another object-oriented language for children that was designed to support the programming of animation [9].

During the 1970s some versions of Logo were created to support the programming of music, color graphics, three-dimensional graphics, and animations. Implementations of Logo appeared on computers inexpensive enough for schools to acquire and the use of Logo by students expanded beyond the laboratory by the end of the decade.

Researchers on intelligent tutoring systems made substantial progress this decade. A notable example is Buggy [10], which was able to diagnosis students' arithmetic mistakes and respond appropriately.

Research on the use of computer games for learning began in this decade as well. Games were developed for educational purposes and researchers explored the educational value of games designed for entertainment purposes [11][12]. The first computer game, MIT Space War, created in 1961, attempted to have accurate positioning of stars and simulation of gravity and hence could be argued to be "educational". Seymour Papert later argued that more serious learning can result from challenging entertainment games than with many "edutainment" games that attempt to be both educational and entertaining [13]. Educational games and educational uses of commercial games has continued to be an active area of development and research for nearly fifty years.

3. 1980s

With the spread of relatively inexpensive personal computers, programming languages for children became widespread in schools and the home. This, combined with Seymour Papert's very influential 1980 book, Mindstorms: Children, Computers, and Powerful Ideas, led to an explosion of activities around Logo. Many schools in the US required its teaching. It became part of the UK National Curriculum in 1988. Far too often, however, the spirit of Logo was lost and children were taught Logo in a way that was far from the creative, exploratory, reflective style it was designed for.

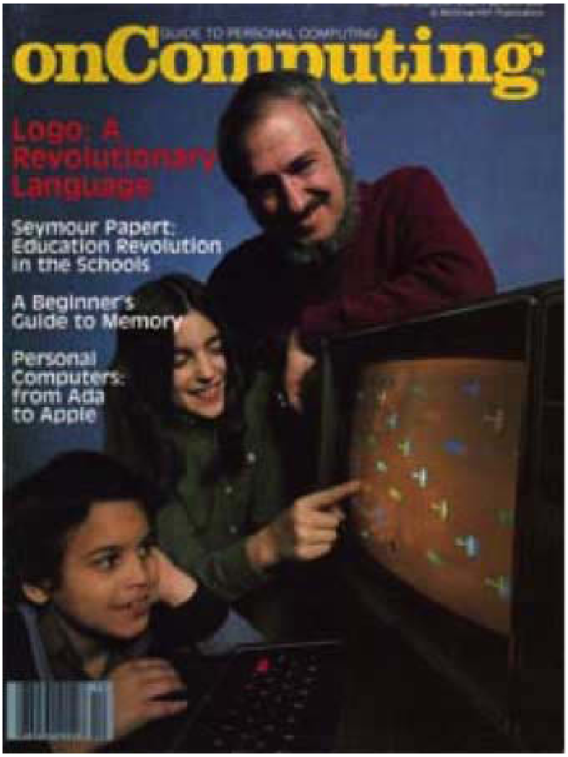

Figure 3. Logo becomes mainstream

Abelson and diSessa wrote Turtle Geometry: The Computer as a Medium for Exploring Mathematics, a book that explores how advanced mathematics could be taught building upon the turtle geometry of Logo [14]. While this undoubtedly helped counter the misconception that Logo was only for primary school children, it was commonly held that Logo was too childish for use by older students. A three-volume book, Computer Science Logo Style by Brian Harvey [15], was aimed at high school teaching and was partially successful in countering this. This misconception is particularly ironic given that Logo was based upon Lisp, an AI programming language with very powerful primitives for dealing with symbolic information.

This decade saw a flourishing of experimental variants of Logo and other programming languages for children. Object Logo [16] was an object-oriented programming language that contained classical Logo as a sub-language. Multi-Logo [17] explored Logo running in multiple processes. Boxer [18] tightly integrated a powerful Logo dialect with a sophisticated user interface. Efforts were made to take other artificial intelligence languages and adapt them for use by school children [19] [20].

Intelligent tutoring systems made strong advances but only in a few select subjects such as teaching algebra, geometry, or computer programming [21].

4. 1990s

The 1990s saw a good deal of activity around adding concurrency and visual syntaxes to programming languages for children. One of the drivers towards concurrency was agent-based modeling. The idea is that one can learn about complex systems by constructing, observing, and experimenting with simulations of interacting entities. This began with StarLogo [22] to be followed by NetLogo [23] and Agentsheets [24]. These efforts to introduce agent-based modeling to school children were described in Mitchel Resnick's book Turtles, Termites, and Traffic Jams [22]. By using these tools, students could acquire a deeper understanding of the underlying processes in scientific phenomena. Topics include those in the physical, biological, and social sciences as well as the humanities including history, philosophy, and language. The educational value of computer programming expanded by providing new ways of learning most school subjects.

Concurrency appeared in other programming languages for children. Stagecast Creator [25] was based upon concurrent rewrite rules. ToonTalk [26] followed the design philosophy of Logo to child-engineer the best computer science programming language ideas. Three decades after Logo's design borrowed from Lisp, ToonTalk's design built upon the ideas of concurrent constraint programming [27]. All of these languages supported programs with multiple simultaneous activities, but only ToonTalk provides general mechanisms for communication and coordination between multiple processes.

The other major trend in the 1990s was to explore graphical syntaxes for programming languages. Agentsheets and Stagecast Creator (then called KidSim) supported expressing programs as graphical rewrite rules. For example, here is how one expressed in KidSim that a character should jump over obstacles [25]:

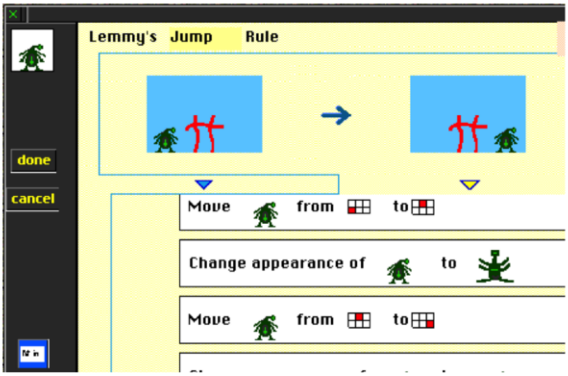

Figure 4. A KidSim rule for jumping over fences

These graphical rewrite rules are intuitive and surprisingly expressive but support abstract rules poorly. Agentsheets addresses this by combining graphical rewrite rules with a spreadsheet metaphor and a scripting language for advanced users.

Agentsheets and Stagecast Creator/KidSim also supported program construction by demonstration. ToonTalk took this to the extreme: the only way to construct programs was via demonstration followed by removal of details to obtain abstraction [28]. ToonTalk has no static syntax; programs are created and viewed as animations in a game-like environment. Programming by demonstration in ToonTalk can be successfully performed by preschoolers [29]. The lack of a static syntax does interfere with scanning and editing programs however. Unlike other programming languages for children, ToonTalk programs can be completely text-free, making them particularly suitable for pre-literate children and internationalization.

In 1996 LogoBlocks [30] pioneered a graphical syntax that subsequently became hugely popular. It introduced shaped blocks that can be dragged and dropped to assemble programs. These blocks correspond to program commands, expressions, data, and control structures. Palettes of blocks enable users to construct programs by selecting the needed parts. Most importantly, these blocks snapped together only when the parts fit together like a jigsaw puzzle. Syntax mistakes are not expressible in such a system. Unlike textual programming languages, the user doesn't need to remember what primitives are available, but instead can select them from palettes. In the next decade the syntactic ideas of LogoBlocks were integrated with StarLogo TNG [31], Scratch [32], Snap! [33], MIT App Inventor [34] and many more [35].

Another programming language trend of the 1990s was to support robot construction kits. A pioneering example of this was LEGO/Logo [36]. As Seymour Papert wrote [37], "LEGO/Logo is a computer-based system that offers a new approach to elementary science education. LEGO/Logo places engineering and design activities at the center of the science curriculum. Using the system, students build machines out of LEGO building pieces (including gears, motors and sensors), connect the machines to a computer, [and] then write computer programs to control the machines. These activities can provide a more meaningful and motivating context for learning traditional science-curriculum concepts while also introducing elementary school students to important engineering and design concepts that are rarely addressed in today's curricula."

Idit Harel [38] and later Yasmin Kafai [39] explored the idea of children programming educational games for younger children. Children using the Logo programming language designed and implemented games to teach concepts about fractions to younger children. Kafai created a sustainable school culture consisting of three grade levels. The oldest children built the games for the youngest children with assistance from middle children who the next year became the game makers. Studies demonstrated that the children who designed and constructed educational games learned the subject matter of their games very well even if the games themselves were not particularly pedagogically effective for the younger students.

The 1990s also saw the rise of multi-media CD-ROMs. For example, Microsoft's Encarta encyclopedia included much that paper alternatives lacked, including audio, animations, videos, and interactive applications. Subjects could be connected by hyperlinks. Novel interactive books on CDROMs where illustrations were animated and reacted to clicks became very popular [40]. So-called "edutainment" games appeared on many CD-ROMs in the 90s.

5. 2000s

The most interesting developments in the first decade of the 21st century were creative and game-changing uses of the Internet. An early example of this was the European WebLabs project [41] which supported children in exploring mathematics and science computationally and sharing and discussing their discoveries in web reports. This added extra dimensions to their learning. In publishing on the web students reflected deeply about what they discovered and worked hard to communicate it effectively. The discussions attached to each report often contained constructive criticism and suggestions. The students were not only doing science and exploring mathematics by constructing computer programs, but were also engaged in the process of academic publication to peers.

The Modelling4All project [42] built a web-based tool (the Behaviour Composer) to support teaching, research, and public engagement with agent-based modelling (ABM). By building on the popular open-source NetLogo agent-based modelling system, the project was able to focus upon higher-level issues of enabling a range of users, including those with no programming experience, to produce open, modular, transparent, sharable models. The Behaviour Composer is web-based both in the sense that one can construct and run models from a modern web browser as well as in supporting sharing models, model components, and interactive tutorials as public web pages.

The Scratch programming language from MIT became very popular after launching its website in 2007. Seven years later, over 7 million projects have been shared on the website, over 4.5 million users registered, and 35 million comments posted. Users support and learn from each other. About 30% of projects are "remixes" where someone makes a variant of another's project (with attribution maintained) [43]. As discussed earlier, Scratch's syntax contributes significantly to its popularity; however, the popularity of other child-engineered programming environments with a similar syntax pales in comparison. The website provides support, motivation, millions of sample projects, and a sense of community that accounts for the popularity of Scratch.

Second Life, a communal three-dimensional virtual world, became very popular in the 2000s. Thousands of avatars controlled by their "owners" interact in this virtual world. "Residents" of Second Life can earn virtual money, build virtual objects, buildings and spaces, and communicate with other residents. A teen-only Teen Second Life was launched in 2005. Educators saw this as potentially a new and effective place for teaching and learning. Many museums opened up Second Life "branches" that exploited the unique capabilities of this virtual world. For example, the US Air and Space Museum built replicas of rockets that visitors could enter and launch. Schools and universities also opened locations. Some uses were recreations of ordinary lecture- oriented teaching while others explored new possibilities. For example, Dr. Peter Yellowlees created virtual hallucinations based upon the experiences of schizophrenia patients. Visitors could experience first-hand what it's like to have schizophrenia [44].

Due to the appearance of inexpensive micro-controllers, educational robotics kits evolved from being cabled to a controlling personal computer to running programs inside the robot itself. Lego's Mindstorms (inspired by Seymour Papert's book of the same name from 1980) became popular. Robot behaviors were still programmed on personal computers, but once downloaded into a micro-controller, they became autonomous. Students used these kits to make a wide range of interactive gadgets. The Lego Corporation offered RoboLab, a graphical dataflow language, to schools using Mindstorms. Researchers implemented dozens of other languages for controlling Mindstorms bricks.

2006 saw the launch of the One Laptop per Child project [45]. The dream was to support the dissemination of inexpensive laptops to every child in the developing world. Special hardware and software was developed. The laptops were designed to have very low power requirements so that electricity could be provided by other means if electrical power wasn't available. The laptops can easily be connected in a network to share

Figure 5. The US Air and Space Museum in Second Life

resources and support multi-user applications. Over two million laptops were produced and in a few countries there were enough to provide a laptop to each child (Uruguay for example). Two million is a significant number, but much fewer than the hundreds of millions initially expected.

6. 2010s

By 2010 web-based technology (JavaScript, CSS, and HTML5) began to be mature enough that serious programming environments could be built to run in any modern browser, including those on tablets and smartphones. Implementations of Logo, ToonTalk, and dialects of Scratch appeared that ran immediately in a browser without any installation or plugins. Programs could be stored seamlessly to cloud storage so that students could move easily between school, home, and libraries as they constructed computational artefacts.

Snap! is a more powerful variant of Scratch, implemented as a web application [33]. It contains new primitives for supporting first-class functions (functions that can create or use other functions) and lists. Unlike Scratch, it is suitable for an advanced high school or beginning university computer science course. It illustrates a tension between programming languages designed to be easy to learn, such as Scratch, and those designed to support more advanced computational concepts and the construction of larger, more complex programs.

ToonTalk was built as a Microsoft Windows application. ToonTalk Reborn is a reimplementation and redesign for the web [46]. ToonTalk programs can be associated with any browser element, giving them interactivity. Widgets constructed in ToonTalk can be embedded inside web pages. Programs and widgets can be dragged between browsers. Programs can be published as automatically generated web pages surrounded by editable rich text.

The Khan Academy [47] began as an online mathematics learning site relying heavily on short videos. It has delivered over 400 million lessons in many school subjects. It delivers 4 million exercise problems daily. Many teachers use it to "flip the classroom" where watching videos as homework replaces classroom lectures. This frees up classroom time for personal support of students as they attempt to do exercises.

This decade has also seen the rise in online tutorials and puzzles designed to teach programming. Code.org has promoted the "Hour of Code" which has reached almost 50 million people. A very impressive online programming tutorial is from the Khan Academy [48]. Each programming lesson replays the actions of an expert with audio commentary. The web page is split between the coding area and an area displaying the result of running the code. Edits of the code are immediately reflected in the output/visualization area. Students can at any time pause the playback and experiment with their own edits or additions to the code area and receive instant feedback.

Another trend of this decade is the programming of smart phones. The MIT App Inventor [34] enables learners to build Android apps in a web browser that can be run either on a phone or in a phone emulator in the browser. It relies upon a variant of the block syntax made popular by Scratch. Pocket Code [49] enables learners to build phone apps on their phones. Its block syntax and interface was designed to work on small screens and touch sensitive devices.

Massive open online courses (MOOCs) became a hot topic in computer-supported learning when in 2011 Stanford University offered a free course Introduction to AI. Its enrolment quickly reached 160,000. Since then courses have been offered at hundreds of universities world-wide with total enrolment of many millions [51] [50]. Because of the large numbers of students, MOOCs generate "big data". This data can be mined to continually improve courses based upon solid evidence.

7. Looking back and forward

’The best way to predict the future is to invent it.’ Alan Kay [52].

Great inventions in using computers to supporting learning have been made during the last fifty years. These include programming languages designed for children, intelligent tutoring systems, online courses, shared virtual spaces, robotics kits, and thousands of games. And learning doesn’t stop with software specifically designed for education but includes use of mainstream developments such as Wikipedia, Google Earth and Maps, social media, computer graphics and animation authoring systems, photo and video editing, 3D printing, spreadsheets, presentation tools, and collaborative document editors.

As computational technology becomes widespread and matures and as the price of computational hardware decreases, we get closer to the fulfilment of the dreams of Seymour Papert, Nicholas Negroponte, and many others that learning by every child on the planet can change dramatically for the better. Children increasingly have devices that enable them to creatively express themselves in a medium that brings their ideas and creations to life. In doing so they acquire powerful ideas for becoming better problem solvers, thinkers, and learners.

References

[1] M. Greenberger, editor, Computers in the World of the Future, MIT Press, Cambridge, MA, 1962.

[2] Dartmouth College Computation Center, A Manual for BASIC, the elementary algebraic language designed for use with the Dartmouth Time Sharing System. Archived from the original on 2012-07-16. http://www.bitsavers.org/pdf/dartmouth/BASIC_Oct64.pdf, 1964
[3] Seymour Papert, Mindstorms: Children, Computers, and Powerful Ideas, Basic Books, 1980.
[4] http://www.platohistory.org/blog/timeline/
[5] Larry Hardesty, "The MIT roots of Google's new software", MIT News Office, August 19, 2010, http://newsoffice.mit.edu/2010/android-abelson-0819
[6] Gerald Sussman, HACKER: A model of skill acquisition, MIT doctoral thesis, 1973.
[7] Ira Goldstein, Understanding simple picture programs, MIT doctoral thesis, 1974.
[8] Leonel Morgado, Maria Cruz, and Ken Kahn, "Radia Perlman A pioneer of young children computer-programming", Current developments in technology-assisted education, Proceedings of m-ICTE, 2006.
[9] Ken Kahn, "Director Guide", Technical Report 482B, MIT AI Lab, December 1979.
[10] John Seely Brown and Kurt VanLehn, "Repair Theory: A Generative Theory of Bugs in Procedural Skills", Cognitive Science, vol. 4, no. 4, pp. 379-426, 1980.
[11] Barbara White, Designing Computer Games to Facilitate Learning, MIT doctoral thesis, 1981.
[12] Thomas Malone, "Toward a theory of intrinsically motivating instruction." Cognitive science 5.4 (1981): 333-369.
[13] Seymour Papert, "Does Easy Do It? Children, Games, and Learning", Game Developer, June 1998.
[14] Harold Abelson and Andrea diSessa, Turtle Geometry: The Computer as a Medium for Exploring Mathematics, MIT Press, 1981.
[15] Brian Harvey, Computer Science Logo Style, Volumes 1, 2 and 3, MIT Press, Second edition, 1997.
[16] Gary L. Drescher, "Genetic AI: Translating Piaget into LISP", Instructional Science, Volume 14, Issue 3-4, pp 357-380, May 1986.
[17] Mitchel Resnick, "MultiLogo: A Study of Children and Concurrent Programming", Interactive Learning Environments, 1:3, 153-170, 1990, DOI: 10.1080/104948290010301
[18] Andrea A. diSessa, "Twenty reasons why you should use Boxer (instead of Logo)", M. Turcsányi- Szabó (Ed.), Learning and Exploring with Logo: Proceedings of the Sixth European Logo Conference. Budapest Hungary, 7-27, 1997.
[19] Ken Kahn, "A grammar kit in Prolog", M. Yazdani, editor, New Horizons in Educational Computing. Ellis Horwood Ltd., 1984. Also In Instructional Science and Proceedings of the AISB Easter Conference on AI and Education, Exeter, England, April 1983.
[20] Richard Ennal, "Teaching logic as a computer language in schools," Proceedings of the First International Logic Programming Conference, Marseille, France, September 1982.
[21] John R. Anderson, C. Franklin Boyle, Albert T. Corbett, and Matthew W. Lewis, "Cognitive Modelling and Intelligent Tutoring", Artificial Intelligence, Vol 42, pages 7-49, 1990.
[22] Mitchel Resnick, Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds, MIT Press, Cambridge, MA, 1994.
[23] Uri Wilensky, NetLogo, http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL, 1999.
[24] Alex Repenning, "Creating User Interfaces with Agentsheets," 1991 Symposium on Applied Computing, Kansas City, MO, IEEE Computer Society Press, Los Alamitos, pp. 190-196, 1991.
[25] David C. Smith, Allen Cypher, and James Spohrer, "KidSim: Programming Agents without a Programming Language", Communications of the ACM, 37(7, pp. 54 - 67), July 1994.
[26] Ken Kahn, "ToonTalk – An Animated Programming Environment for Children", Proceedings of the National Educational Computing Conference, Baltimore, Maryland, June 1995. Extended version in the Journal of Visual Languages and Computing, June 1996.
[27] Vijay Saraswat, Concurrent Constraint Programming, MIT Press, Cambridge, MA, 1993.
[28] Ken Kahn, "Generalizing by Removing Detail", Communications of the ACM, 43(3), March 2000. Extended version in Henry Lieberman, editor, Your Wish Is My Command: Programming by Example, Morgan Kaufmann, 2001.
[29] Leonel Morgado, Maria Cruz, and Ken Kahn, "Working in ToonTalk with 4-and 5-year olds", Proceedings of the IADIS International Conference, e-Society, Vol. II, IADIS, 2003.
[30] Andrew Begel, LogoBlocks: A Graphical Programming Language for Interacting with the World, MIT Advanced Undergraduate Project, http://research.microsoft.com/en-us/um/people/abegel/mit/begel-aup.pdf, May 1996.
[31] Eric Klopfer, Hal Scheintaub, Wendy Huang, Danial Wendel, and Ricarose Roque, "The Simulation

Cycle: combining games, simulations, engineering and science using StarLogo TNG", E-Learning and Digital Media, Volume 6 Number 1, 2009.
[32] Mitchel Resnick, John Maloney, Andrés Monroy-Hernández, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, and Yasmin Kafai, "Scratch: Programming for All", Communications of the ACM, Vol. 52 No. 11, Pages 60-67, 2009.
[33] Brian Harvey and Jens Mönig, "Bringing 'No Ceiling' to Scratch: Can One Language Serve Kids and Computer Scientists?", Constructionism 2010 Proceedings, Paris, 2010.
[34] David Wolber, Hal Abelson, Ellen Spertus, and Liz Looney, App Inventor, O'Reilly Press, 2011.
[35] Blockly, https://developers.google.com/blockly/
[36] Mitchel Resnick, Stephen Ocko, and Seymour Papert, "LEGO, Logo, and Design" Children's Environments Quarterly, vol. 5, no. 4, 1988.
[37] The Laboratory Schools LEGO-LOGO Project, http://www.ucls.uchicago.edu/students/projects/1994-95/Lego-Logo/ProjectDescription.html
[38] Idit Harel, Children Designers: Interdisciplinary Constructions for Learning and Knowing Mathematics in a Computer-rich School, Ablex Publishing, 1991.
[39] Yasmin B. Kafai, Minds in Play: Computer Game Design as a Context for Children's Learning, Mahwah, NJ, Lawrence Erlbaum, 1995.
[40] Wikipedia, "Living Books series", https://en.wikipedia.org/wiki/Living_Books_series
[41] Yishay Mor, Richard Noss, Ken Kahn, Celia Hoyles, and Gordon Simpson, "Thinking in Progress", Micromath, The Association of Teachers of Mathematics, 20(2), pp.17-23, 2004.
[42] Ken Kahn and Howard Noble, "The Modelling4All Project – A web-based modelling tool embedded in Web 2.0", Proceedings of Constructionism 2010, Paris, August 2010.
[43] MIT Media Lab, "Scratch statistics", http://scratch.mit.edu/statistics/
[44] BBC News, "What it's like to have schizophrenia", http://news.bbc.co.uk/1/hi/health/6453241.stm
[45] One Laptop per Child, http://one.laptop.org/
[46] Ken Kahn, "ToonTalk Reborn, Re-implementing and re-conceptualising ToonTalk for the Web", Proceedings of Constructionism 2014, Vienna, August 2014.
[47] Khan Academy, https://www.khanacademy.org/
[48] Khan Academy, "Computer programming", https://www.khanacademy.org/computing/computer- programming
[49] Wolfgang Slany, "Tinkering with Pocket Code, a Scratch-like programming app for your smartphone", Proceedings of Constructionism 2014, Vienna, August 2014.
[50] Luc Steels (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam. 2015.
[51] Wikipedia, "Massive Open Online Course", https://en.wikipedia.org/wiki/Massive_open_online_course
[52] Wikipedia, "Alan Kay", https://en.wikipedia.org/wiki/Alan_Kay

Music Learning with Massive Open Online Courses (MOOCs)
L. Steels (Ed.)

doi:10.3233/978-1-61499-593-7-225

Chapter 15. Blended Learning And MOOCs

Georges Van der Perre
K.Universiteit Leuven

Abstract. The digital transformation of the world is likely to cause dramatic shifts in the world of learning, as is shown spectacularly by the MOOCs phenomenon [5]. It faces universities with the opportunity (and the obligation) to transform themselves thoroughly into institutions with a significantly improved and extended service to society and the ability to adapt flexibly to the rapidly changing needs (learning organizations). The Internet and the multimedia-interactive information technology allow to extend education beyond the (school) hours as well as beyond the (class) walls. The consequences of this simple observation are not yet fully seen. It is not just that universities' own regular full-time students get new "blended" learning schemes offered, but it also means new tools for flexible, part-time and distance learning, and it especially implies that external target groups can be served significantly better, such as future students (study orientation), graduates and professionals (lifelong learning), and last but not least international (global) audiences (virtual mobility).

A case study initiated in 2014 by the Royal Flemish Academy of Belgium for Sciences and Arts and lead by two prominent thinkers from abroad concentrated on the question how traditional universities should deal with blended learning and MOOCs. The key conclusion is: "The optimal exploitation of ICT and the Internet for the new higher education of the 21st century" will not take place spontaneously. A 'bottom up' approach, i.e. the support of a multitude of individual initiatives, is necessary to let creative ideas grow, but insufficient to bring about the necessary changes in higher education. This requires powerful and radical 'top down' measures, and some concrete recommendations are given in this respect. In addition there is a continuing need for further fundamental research and visionary thinking.

Keywords. Blended learning, online learning, MOOCs, universities, higher education, distance learning, flexible learning, lifelong learning, virtual mobility, digital technology.

1. A systemic vision as starting point.

The world today is completely permeated by digital technology propagating deeper and wider every day. The university campus can be seen as a physical artefact built on top of a digital entity [2]. The campus houses as much in digital data, information and communication systems and software as it houses in buildings and premises. Unlike the buildings however digital technology creates an opening to the world outside the campus: companies, society, the other universities here and around the world, anyone who owns a PC, laptop or smart phone. This opening is manifested in a very spectacular way by the MOOCs.

Universities have not yet really adapted their strategy to this digital reality. This is not just about online learning and blended learning. For example, from the huge amounts of data (big data) that universities can collect about their students today and tomorrow, it becomes possible to design digitally an optimal curriculum for each student individually, including customized learning methods. These data provide at the same time a huge source of information for the development of educational strategies and methods of the university, the faculty and the teachers themselves. Obviously digitalization has enormous implications for scientific research as well, but those are beyond the scope of this project.

Students and teachers living on and around the campus today have grown up in this digital world and want to develop further in it. The practice of higher education lags behind on this reality, and this gap seems to be growing yet. There is an increasing alienation of the learner from the currently used methods of learning. The development of digital technology is the at basis of dramatic shifts in the world of learning, whether we like it or not. At the same time universities are faced with the opportunity (and the obligation?) to transform themselves thoroughly, into institutions with a significantly improved and extended service to society and the ability to adapt flexibly to the rapidly changing needs (learning organizations).

Universities would not be universities if they would not handle this transition with care and criticism, and have no regard for the other side of the coin: the potential adverse effects of the use of learning technology on learning (particularly ”profound learning”), on the personal development of students, and on the university itself.

This paper is the outcome of a case study organized in 2014 by the Royal Flemish Academy of Belgium for Science and the Arts (KVAB) entitled ‘Blended Learning in Flemish Higher Education’. Similar studies have appeared in other European countries. For this project two thinkers from abroad were invited to explore how online learning resources could be blended with more traditional forms of university education. Both Thinkers have complementary top expertise in the advanced use of ICT in education: Diana Laurillard (University College London, London Knowledge Lab) is a leading academic in the area of blended learning and Pierre Dillenbourg(Ecole Polytechnique F’ed’erale de Lausanne,EPFL) is the coordinator of the comprehensive MOOCs program at the EPFL.

For a year, the Thinkers collaborated closely with a representative local expert group. They also participated in various seminars and workshops at the five Flemish Universities and the UCL (Universit’e Catholique de Louvain). During a closing symposium on November 19, 2014 they presented the results of their study, and confronted them with the views of representative organizations and the experiences of providers and users of online courses. The Thinkers have cast their opinions and recommendations into two position papers that are published integrally in the KVAB- Standpunt 33 report [4],[2]. A synthesis of the position papers appeared in the journal TH&MA [6].

The paper summarizes the main findings of this case study. It discusses the new opportunities generated by online learning, how to unlock and use them, and what strategies universities can use. The paper concludes with some personal observations based on my own long term involved with long distance education.

2. Opportunities

2.1. New outward opportunities

Widening access in time and space.

The Internet and the multimedia-interactive information technology take education (or at least some forms of it) outside the (school) hours as well as outside the (class) walls. This statement is very simple, but nevertheless its consequences are not yet fully seen. It is not just that universities' own regular full-time students get new learning schemes offered (such as flipped class, see below), but it also means new tools for flexible and part-time learning, and it especially implies that external target groups can be served better, such as future students (study orientation) and graduates (lifelong learning).

From classes to masses

In the digital world almost everyone has almost unlimited access to a vast global source of information (libraries, databases, Wikipedia ...) .But 'a library is not yet a university'. Even when MIT launched its Open Course Ware (OCW) project early this century, it was still clearly said: this is course material (primarily for teachers), this is not the teaching. In contrast, the MOOCs resolutely take the step to the actual teaching, all components included: not just video lectures but also problem solving, tests, discussion groups, assignments, exams. They perform a truly spectacular scale leap (both in student numbers and action radius): from "classes" of at most a few hundred (and preferably much less) students on campus to "masses" of one hundred thousand students all over the world. Of course 'it is not all gold that glitters' ("In general good MOOCs are better than bad MOOCs!" [2], and of course the student drop out is generally quite high (around 80

MOOCs as networks for knowledge building

Today's MOOCs are not just media shows of big star professors. In these global networks of teachers and learners also collaborative knowledge building takes place. Sometimes learners become teacher or researcher (peer learning, peer assessment, crowd sourcing) and teachers and researchers become learners. A hundred thousand participants allow a hundred thousand observations, whether on their own learning, on their Italian translation of some English sentences, or about the weather outside. This way observation and measurement data sets are obtained with an exceptionally high statistical significance and reliability.

Since a number of years, scientific research centers have established their own digital networks, the interaction between them is continuous and no longer confined to papers and meetings. MOOCs can become a large scale and open variant of these research networks. Perhaps the potential of a new type of global virtual knowledge centers has just been unveiled, perhaps there are certain scientific, technical, social, medical questions that can find a better and faster response through global interaction among peers than by specialized research in competitive research groups.

2.2. New inward opportunities: blended learning.

Teachers and students today have to their disposal a lot more and better tools for teaching and learning than this was the case at the beginning of this century, both within the classroom and outside (online). Online learning adds a new dimension to education. Students can access learning materials at home and after school hours, view lectures, participate in seminars and discussions, perform tests and get feedback on them, work on projects in group. The decoupling between learning and a fixed time schedule alleviates several logistical problems and creates a lot of flexibility, it allows to provide a full educational support for part-time students, working students, students with special programs. Blended learning is undoubtedly the model of the future, with 'blends' tailored to the specific needs and context of the target group.

With online learning, some well-known problems in current higher education can be solved. Students in the final year of secondary education can be better informed and oriented in their study. Another problem is that of the students who come inadequately prepared to problem sessions and seminars: they can bring their knowledge up to date through online tests with feedback.

Some pioneers of educational innovation apply completely new schemes of teaching. They let their student teams prepare themselves their course in wiki format, using "Open Educational Resources' (OER) available on the web, and give them projects to carry out using that course. They do not give lectures in the classical sense, but animate every week an intensive live session, in which they give presentations and explanations where necessary. For the assessment of students they look not only at the end result (the project and the oral exam about it) but also monitor their activity during the live sessions and on the blogs, and integrate some peer assessment.

But not everything that happens is well thought through. In the flipped class concept students can watch in advance at home video clips in which the teacher presents short modules of the course, before attending interactive lecture sessions with the teacher and his assistants. The model is sometimes applied to solve the problem of the student going unprepared to working sessions , but it is not catching on just like that.

With educational technology a lot is possible, but that is far from saying that it all goes by itself. Technology does not solve the problem itself. And after all, the ancient proverb "what's the benefit of candle and glasses when the owl does not want to see" (again a literal translation of an ancient Flemish proverb: 'wat baten kaars en bril als de uil niet zien en wil') remains valid even for digital learners. Scenarios of blended learning have to be thoroughly thought through, planned, evaluated and adjusted. That brings us to the topic of 'learning design' [4], a new professional and (applied-) scientific discipline.

3. Learning design

Learning design is a design discipline, of the same kind as mechanical design and object-oriented design. One does not just cobble a machine or a software package together, there is some professional knowledge and methodology involved.

Professionalization

Developing a blended learning scenario requires a much more professional approach than setting up a traditional course. Obviously you start from your learning objectives, the needs and potential of your target audience, and the concrete context in which teaching and learning take place. There are different ways of learning: listening or reading ("acquisition"), discovery (research), discussion, practice, group work, assignments. Evaluation (formative and summative) is always an essential component. Conventional teaching is based upon co-presence in time and space and the use of physical objects. Digital technology adds to that online communication and virtual objects. It allows in principle more personalization, flexibility, accessibility, inclusiveness (for students with learning disabilities, disadvantaged groups) and efficiency (also in terms of cost). Whether or not all that potential is also exploited effectively is a matter of ... .learning design.

More work for the teacher?

According to Diana Laurillard teachers who want to start with blended learning, more specifically with applications of online learning, should be aware of a substantial increase in their workload. The list of tasks to be fulfilled in the design and implementation of a blended learning course is very long [4]. This statement might discourage teachers to take the step, and in fact it is questioned by the experts group. There is agreement on the fact that there will be substantive shifts, such as less time for lectures and more time for coaching of small groups. There may have to be cut in other, less efficient tasks of the teaching team. The (formative) evaluation of students can be improved using technology (tests with automatic scoring and feedback already now, learning analytics in the future) , for the exams (summative evaluation) this is not yet generally accepted.

Personal satisfaction: a deeper understanding of the own learning goals

An interesting experience is that designing e-learning applications often leads to a deeper understanding of the intended learning objectives (and even of the essential content of the course matter) by the teacher teams. Those who went through this personal experience sometimes wonder what they were doing before.

How to make this happen?

It seems appropriate that teachers are supported by teams of specialists in learning design, and that at the same time they build among themselves ”communities of learning design” within the various scientific disciplines. Moreover learning design deserves recognition as a full-fledged interdisciplinary scientific discipline, such that also non-educationalists and non-pedagogues can build a full academic career concentrating their research efforts on this subject. Diana Laurillard [4] gives a number of concrete suggestions and recommendations in this respect. In order to make this ”learning design” ”evidence-based”, the sharing of experiences within the ”communities of learning design’ is absolutely vital. But there is not only a need for experience based expertise, also systematical, quantitative, fundamental research is necessary.

The big unknown: learning!

However professionally we design and implement our blended learning scenarios (learning design) , still it often remains a question whether and to what extent we have stepped up the quality of teaching and enhanced the attainment of learning objectives. Two questions for illustration.

- Independent learning (autonomous learning, personal learning) is widely professed in higher education as a point of belief, especially in a constructivist approach. ”Guided personal learning” or similar mottos were on the education banners of universities at the end of the last century. Is this still the case? For some technological applications that enhance the student’s learning comfort it is far from certain that they promote independent learning. Yet, learning technology comes best into its own in a context of independent learning. Open Educational Resources (OER) , Open Course Ware (OCW), learning in virtual groups and networks, online learning per se, are not really booming within mainstream education, precisely because education does not resolutely opt for independent learning, except for some brilliant pioneers (see above).

- Secondly: does online learning lead to deeper learning or is the opposite true? The well-known technology critic Nicolas Carr [1] has serious doubts. The core of his argument (supported by scientific research) is that the simultaneous use of multiple information streams (or even just the possibility to do this, in the case of hypertext) can have a detrimental effect on people's acquisition of deep and durable knowledge. The enormous flow of concepts and data, the fact that they come in all together through multiple senses, and the fact that part of our mental energy is spent on controlling these flows (to click a link or not?) overload our ability for mental processing and prevent that a sufficient fraction of the incoming concepts can settle definitively and correctly in our long-term memory.
  Moreover - and this would be even worse - Carr argues that by this way of learning we slowly but fatally lose the skill to acquire lasting knowledge (deep learning) through deep reading; the brain is plastic and adapts to the most frequent activity, at the expense of the less frequent ones. Carr's ideas are highly controversial and probably strongly overstated, but they raise important questions that require further research.

So there still is a lot of work to do for scientific research. We should not forget that for youngsters to develop themselves with knowledge, skills and attitudes to full-fledged bachelors and masters in a specific discipline is indeed quite different from updating their knowledge in a later age (e.g. with MOOCs.)

4. A systemic approach to innovation in higher education.

The MOOCs are a phenomenon that has dropped out from the sky and spread via the worldwide web, basically unhindered by existing traditions and rules. To develop blended learning within higher education is quite a different story: it is an innovation to be integrated in an existing system with a well-defined societal function; existing traditions, role distributions, expectation patterns; established rules and organizational, managerial and funding models. The complexity of the existing higher education system with its different decision levels, stakeholders, actors and pressure groups makes it highly resistant to change. [4].

Implementation steps

Therefore the implementation of innovations in higher education requires a systemic approach in the following sense:

1. We should start from the fundamental role of higher education, which according to former Stanford president Hennessy consists of learning and accreditation [3]. Learning: to guide every student through a process by which he/she optimally develops his / her potential to her/his own benefit and that of society. During this process he/she acquires knowledge, skills and attitudes which he / she would not be able (or have serious difficulty) to acquire on her/his own [4]. Accreditation: a diploma system should provide the necessary guarantees to society and give the individual legitimate recognition and optimal development opportunities.
2. We should act intelligently upon the ”drivers” that steer the actions of the actors and the ”enablers” that support these actions. Which are these ”drivers’, which are these ”enablers”? To which extent the drivers are innovation-oriented and the enablers innovation-friendly?

- Drivers are: funding systems, rules regarding exams and certification, needs of stakeholders, quality assurance, strategic plans of institutions, curriculum contents requirements, needs and skills of students, careers of the teachers (objectives and opportunities). Many teachers would add spontaneously: the passion of the professor to share his expertise with the young (and the personal reputation within the student audience!)
- Enablers are: leadership and encouragement of educational innovation by the academic authorities, support for professional development (continuing education) of teachers, communities of practice, systems, tools and support services for learning technology, evidence from research and practice, learning materials for exchange and sharing.

From this analysis, it follows that two systemic actions could make a difference (along with a profound reflection on the role of higher education):

- an adjustment of the drivers making them to encourage the development of new practices of blended learning
- a further development of the enablers to make these new blended learning practices feasible, effective and sustainable.

Both these actions require firm and tangible interventions in the system.

The financial picture

Whatever the teaching method used (from conventional to online and all blends in between), there is always a fixed cost and a variable cost.

- The fixed cost for the design and preparation of a course (contents, learning materials, sources, activities, resources, learning environments) is independent of the size of the student cohort.
- The variable cost for teaching and coaching (tutoring, discussion, advice, counseling, guidance, formative evaluation, grading) is a unit cost per student, and thus increases proportionally to the size of the student cohort.

The use of learning technology in blended learning increases the fixed cost and is therefore more cost effective as the number of students is higher. Yet it is precisely the use of learning technology (such as the globally accessible MOOC-learning platforms) that allows a spectacular scale expansion. . If in addition the variable costs can be reduced by automating some of the related functions, it is possible to achieve a financially viable system. At present, the MOOCs are not yet financially viable, at least not for the supplying universities, because they are offered free or below the real cost , or because the share of the course fees that reaches the university is below the cost for the university. [4]. In the context of blended learning inside the university, a spectacular scale expansion is not an option, unless universities will introduce MOOCs into their regular education (see below), or go to develop and offer courses in networks of several universities (see below). Universities often opt for SPOCs: small private online courses. The fixed cost can be reduced by reusing and annually updating the courses and teaching materials. The reduction of the variable cost is the biggest challenge. The partial automation of tutoring and assessment and the development of forms of peer learning (discussion, peer assessment) are adequate strategies to achieve this.

In all of these developments the personal contact between teacher and student (and among the students themselves) must not be lost. An interesting experience in the service sector (e.g.in banking) is that as more functions are efficiently automated ,the personal contacts between the company and the client develop their own intrinsic value and become more and more appreciated. The same trend can be expected in higher education.

5. When should universities engage in MOOCs?

There are many good reasons why European universities should become actors on the MOOCs scene. We can group these in four clusters [2]:

1. Like it or not, it is happening. Recent data from the platforms Coursera and EDX show that in Flanders 50,000 people enrolled for MOOCs in the last two years. The vast majority of them already got a higher education diploma, and only a good fifteen percent of them reach the finish line. But any way, there is a significant part of advanced education that escapes from the control of universities. Our universities do not control (decide, determine) who offers courses in digital space nor what their students can learn there.
2. Better to be an actor than a spectator. The MOOCs (and all the further digital learning developments that will follow from them) will change the value scales and rankings (such as international university rankings) that prevail in the academic world.
3. MOOCs create new opportunities directly related to their big scale (see above), which might be beneficial for the university's own on campus students as well (see below)
4. The current situation of universities is far from perfect anyway : the quality and efficiency of education is subject to improvement, the study orientation and the success rate of first year students is problematic, the universities can do a better job in facilitating flexible learning and supporting working students, ...

Moreover, with MOOCs the societal role of the university can significantly be strenghtened and widened. Lifelong learning is a broad action field for digital education: ”a diploma with a service contract” for the universities' alumni, training courses for companies and public services, continuing education for teachers, ... Digital learning allows to meet specific target group needs: preparation and orientation of students in the final year of secondary education, retraining programs in view of employment. With ”agile” digital curricula universities can respond quickly to new developments that create new training needs for which the classical curriculum revision procedures are too slow. And finally, the public debate can be enriched by the contribution of complete and reliable information and science-based insights from the academic world.

The encounter between two worlds. MOOCS: the Trojan horse for university education?

Within Flemish universities we found that there are strong reservations against the integration of online courses, and especially online courses of foreign origin such as MOOCs, into their own degree programs. By 'integration of courses' we mean: assigning credits (ECTS credits) to them in the regular study programs. Universities which produce MOOCs themselves (as the nearby universities EPFL Lausanne and UCL Louvain) have opened some of them (and assigned credits to them) to their own on-campus students, with varying success. Here too, some learning design is required, and the EPFL has already built some valuable experience in this area [2].

MOOCs are becoming large scale global communities of practice where novice and mature students, teachers and professional people are building knowledge together. Why shouldn't we give our students access to this new world?

In this respect it is an encouraging signal that the Dutch-Flemish accreditation body NVAO as well as the representative Flemish educational councils VLOR and VLIR expressed their approval of the integration of MOOCs in regular high education study programs . At least as important is the fact that the student body VVS is in favor of it. Its only caveat is that the integration of online learning should not diminish the direct contacts with the teachers and the quality of the tutoring.

Since twenty years there have been efforts to set up inter-university networks for online learning in Europe, at national levels as well as on a European scale. Within these networks universities would exchange and share their courses online and jointly produce and organize online courses (including professional training courses for the industrial world) . Although the pilot experiments were generally successful, most of the European initiatives evaporated once the funding from the EU stopped. And yet, for the inter-university exchange of online courses European universities have a competitive edge with respect to the US and other regions: the Bologna treaty is a basis for cooperation and student mobility, and the ECTS credit system provides the "single currency" for the exchange of courses. But apparently these networks have as far never been seen as a strategic priority by their member institutions: universities are more focused on competition than cooperation. Is this now going to change under the pressure of the MOOCs reality? It does not look like, despite the obvious value of these networks for teachers and (especially) the students ("Virtual Mobility").

Concrete Measures towards blended learning

The key message is: "The optimal exploitation of ICT and the Internet for the new higher education of the 21st century" will not take place spontaneously. A "bottom up" approach, i.e. the support of a multitude of individual initiatives, is necessary to let creative ideas grow, but insufficient to bring about the necessary changes in higher education. This requires powerful and radical "top down" measures:

- The relentless growth of digital technology in our daily lives will undoubtedly continue to affect higher education. But for the optimal exploitation of the potential of ICT for learning more is needed than just a further introduction of technology in an otherwise unchanged teaching and learning system. A systemic, holistic approach is called for, aiming at the transformation of the higher education system in such a way that it allows to deploy the best of what technology has to offer for a thoroughly innovated education and learning methodology.
- Blended learning is an art and a skill which needs and deserves to be raised to a higher level: learning blends must be developed professionally. A specific design skill is called for: learning design. At the same time scientific research is needed to build a fundamental basis for this design skill such as to make it evidence based. We plead for the recognition of learning design as a true domain of interdisciplinary fundamental and applied research. We suggest that each educational institution creates a central interdisciplinary department for educational innovation, in which four functions are to be integrated: fundamental research, development of tools and techniques, training and support for teachers and initiation of concrete innovation projects in collaboration with the teaching staff in the different academic disciplines (faculties and departments).

- On the international MOOC scene we believe that interuniversity collaboration at the Flemish and European level in the form of so-called DOCCs (Distributed Open Collaborative Courses) is the way to go. Interuniversity collaboration not only makes MOOCs financially feasible, it also significantly enriches them in contents and quality. Flemish Higher Education institutions are strongly advised to participate actively in the OpenupEd initiative led by EADTU and financed by the European Union. Through the attribution of credits to MOOCs and DOCCs for students in regular education we give them access to these exciting European and global developments and allow them to interact with professors and peer students all over the world (virtual mobility). With its Bologna Declaration and its ECTS (European Credit Transfer System) Europe has developed both the basic vision and the practical instrument for this virtual mobility.
- For the coordination and the support of the interuniversity collaboration at the Flemish level, as well as the collaboration with companies, professional organizations, government and other stakeholders, a central institute is called for. This institute should be governed by a board consisting of the vice rectors for digital education of the universities as well as delegates from industry, government and societal stakeholders. It is not only to promote and to support (technically and logistically) all forms of collaboration between partners, but it can also play a leading role in research and initiate and manage specific projects such as the fast creation of MOOCs or DOCCs as an agile response to emerging needs of society and industry. The Dutch SURF could be taken as a starting model. Also the existing collaboration between the Flemish universities and Dutch Open University is due for a creative rethink in view of the digital revolution, the convergence of on campus and distance education and the complementary expertise of the classic Flemish universities and the open university.

6. Some Personal conclusions

I now come back to the starting point of the project, which at the end appears to be one of its main conclusions: ‘there is a need for a systemic vision on the optimal exploitation of ICT and the internet for the new learning of the 21st century’. It leaves me with a double question. First: did we go wide and deep enough in developing new insights that can be step stones and building blocks for this ‘systemic vision?’ Second: will the views and recommendations presented above actually bring about tangible progress in ‘the optimal exploitation of ICT for the new learning’?

Let me start with the second question as it strongly affects the answer to the first one. It can in itself again a be split in two parts: a) have the presented views and recommendations a chance to be accepted and implemented by the authorities and people they are intended for, and b) if they are, will that lead to tangible and sustainable results? The first part will first of all depend upon the follow up actions taken after the publication of the project reports, a process which is going on right now. But for both parts we can learn lessons from the last twenty (!) years.

I take the liberty to refer to a chapter I wrote in 2007 [8] under the title ‘A European virtual learning area, now or never?’ The basic issue I discuss in that paper is the following. Most of the models for e learning we are discussing now (i.e.in 2007) were developed and even tested out in large scale pilot experiments (eg. coordinated by EuroPACE) in the nineties. They never really penetrated in mainstream higher education and did not even show sustainable in corporate training. Most of the national virtual or digital universities that were founded around the beginning of this century (e.g. the Digitale Universiteit Nederland, the UK eUniversity etc.) did not survive. So why would it all work now? I quote

> 'In comparison with ten-fifteen years ago, there are some major differences, but some basic limitations remain. A first major factor of difference is technology, ... Academia at large has accommodated the basic (technical) elements of e Learning. A second major factor of difference is 'internationalisation'. International networking is high on the strategic agendas of universities ... Which basic limitations remain? There are first the intrinsic limitations of e learning as such. E learning is basically sitting in front of a PC or a projection screen, ... Secondly there is the basic motivation of 18-25 years old on campus students. Their first motivation is to have a five years 'sabbatical' before going into real professional life, and to make this sabbatical fulfilling and enriching in a number of ways. Their study in the strict sense is only one aspect of their sabbatical, the key objective of their study is the diploma and the key factor the confort and support they get to obtain it. Therefore, replacing lectures, labs and tutorials by effective ways of interactive e learning is not an option for them ... Thirdly there is the strategy of the universities. I had to learn the obvious truth that the very first goal of every university is the same as that of any organisation: to assure and extend its own existence ... So positive new trends and opportunities in learning are only supported wholeheartedly if they serve the interests of the universities and in particular 'this university' ... Nevertheless ... I see two positive trends on the scene of networked e learning in European universities: the 'restricted virtual learning areas' movement and the 'open virtual learning area' movement. In the former, universities carefully select or build their networks for e-learning as a tool for the implementation of their 'internationalisation' strategies . In the latter, universities follow the trend set by MIT to make their e learning courses freely available through the web. The two movements seem contradictory but are not necessarily in competition, as the first is about education and the second about course materials. '

Reading this all now, I realize that the example of a 'restricted virtual area' I had in mind was EUNITE, a network that has been dissolved a couple of years later. With the term 'open virtual learning area' I referred to the Open Educational Resources (OER) and 'Open Courseware' (OCW) movements. Many universities participate in these, which does not necessarily mean that OER are significantly used in their teaching practice. And finally, where I wrote that 'open virtual learning area' initiatives such as MIT's are about course materials and not about education, this is not at all the case anymore since 2011 and the outbreak of the MOOCs tsnunami. There is no doubt that the MOOCs - and the prestigious top universities behind them-have caused a breakthrough in the thinking about eLearning in universities all over the world. Did they bring about a disruptive change in the higher education system? That is quite another question, and many even doubt whether they will survive as such.

The goal of the above made critical reflections is not to succumb to defeatism but to learn from the past. In setting up new initiatives it is useful to understand which are the critical success factors and to see the technological and societal developments that allow us to be successful now in creating things that were not possible in the past. This brings me back to the first question: did we go wide and deep enough in developing new insights that can be step stones and building blocks for this 'systemic vision'? My answer is yes and no. It is yes for the short and intermediate term and the present structure of higher education. We have been very wise to stick to reality and I am confident we developed valuable and useful insights and strategies that will lead to concrete results.

My answer is ‘not yet’ for the longer term. I have the feeling that we are not yet fully aware of the implications of digital technology for knowledge as such and its creation, for global communication, for sharing and jointly and interactively creating knowledge, for the development of personalized learning contents and support, for independent and open learning, for learner support and assessment, for monitoring and adaptive control of learning processes (learning analytics). We still have much to learn about learning as such, e.g. how to stimulate ‘deep learning’. Perhaps we stick to much to the model of today’s universities and today’s education which is still strongly based upon courses and teaching. Perhaps we need a paradigm shift and thorough mental changes among students, teachers, in society at large. How would the university look like if we invented it now? That sounds like a very imaginary question, but it might be a good idea to create such a pilot university. Or would one of the existing universities dare to take the risk of experimenting with a totally new concept?

7. Acknowledgement

The project ‘Blended Learning’(2014) is part of the ‘Thinker in Residence’ program of the KVAB that is carried out in the framework of the covenant between the Royal Flemish Academy of Belgium for Science and the Arts (KVAB) and the Flemish Government.

References

[1] Carr, N. (2010) The Shallows: What the Internet Is Doing to Our Brains. New York: W. W. Norton.
[2] Dillenbourg, Pierre (2014) ’Proposal for a Digital Education Strategy for Flanders Universities.’ In: G. Van der Perre en J. Van Campenhout (eds.) (2014/2015) Higher education in the digital era. A thinking exercise in Flanders. KVAB-Standpunt 33, Brussel.
[3] Hennessy, J.L. (2012) John Henessy: Risk Taker. IEEE Spectrum, May 2012.
[4] Laurillard, D.. ‘Thinking about Blended Learning. A paper for the Thinkers in Residence programme.’ In Higher education in the digital era. A thinking exercise in Flanders. KVAB Standpunt 33, Editors: G. Van der Perre en J. Van Campenhout. Brussel: KVAB- Thinkers in residence program 2014, 2015.
[5] Steels, L. (2015) The coming of (social) MOOCs. Chapter 1. In: In: Steels, L. (2015) (ed.) Music Learning with Massive Open Online Courses (MOOCs). IOS Press, Amsterdam.
[6] Van der Perre, Georges, Jan Van Campenhout, Pierre Dillenbourg, en Diana Laurillard (2015). Geef studenten toegang tot een nieuwe wereld. Blended learning en MOOC’s: een denkoefening in Vlaanderen.’ TH&MA-Tijdschrift voor Hoger Onderwijs en Management (Instondo B.V.) 22, nr. 1 (Januari 2015): 29-37 (in Dutch).
[7] Van der Perre, Georges, Van Campenhout Jan et al. (2015) KVAB-Standpunt 34: Hoger onderwijs voor de digitale eeuw (Higher education for the digital century) KVAB, Brussels.
[8] Van der Perre, G. (2007) ‘A European Virtual Learning Area. Now or never?’. In: A.Boonen and W. Van Petegem (2007) European networking and learning in the future., The EuroPACE approach. Garant, Antwerpen-Apeldoorn.

Subject Index

Author Index

This page intentionally left blank