Unsupervised Filtering of Color Spectra

Reiner Lenz, Mats Österberg
Image Processing Group, Dept. EE, Linköping University,
S-58183 Linköping, Sweden, reiner@isy.liu.se, mats@isy.liu.se
Jouni Hiltunen, Timo Jaaskelainen
Väisälä Laboratory, Dept. Physics, University of Joensuu,
FIN-80101 Joensuu,Finland,jhiltu@fysnext.joensuu.fi,jaaskelainen@joyl.joensuu.fi
Jussi Parkkinen
Dept. Information Technology, Lappeenranta University of Technology,
FIN-53851 Lappeenranta, Finland, parkkinen@luoti.lut.fi

Abstract

We describe how unsupervised neural networks can be used to extract features from databases of reflectance spectra. These databases try to sample color space in a way which reflects the properties of human color perception.

In our construction of neural networks we identify first desirable properties of the network. These properties are then incorporated into an energy function and finally a learning rule is derived using optimization methods to find weight matrices which lead to minimal values of the energy function. We describe several energy functions and the performance of the resulting networks for the databases with the reflectance spectra. The experiments show that the weight matrix for one of the systems is very similar to the eigenvector system whereas the second type of systems tries to rotate the eigenvector system in such a way that the resulting filters partition the spectrum into different bands. We will also show how the additional constraint of positive filter coefficients can be incorporated into the design.

Unsupervised Processing of Sensory Data

The question how animals and humans process sensory data has received increasing attention recently. This is obviously of great biological interest but it could also help in the solution of technical problems since these systems have obviously learned how to deal with this type of data in an efficient way.

Many of these approaches are based on the observation that the sensors (for example the retina or the ear) produce a large amount of raw data. The information contained in this dataset has to be transmitted to the brain over the nerves. Taking into account the limited capacity of the nerves this is only possible since the data is usually highly redundant. Some processing of the data must therefore take place on the receptor level and the type of processing must be data-driven, i.e. it must take into account the properties of the input data space. The data space investigated in this paper is the space of color spectra that are relevant for human color vision. We will use samples from this space and show how systems can be trained to learn efficient processing of the spectral data. Many of these ideas go back to the fifties and sixties [2], and for some recent contributions the reader may consult [1], [5], [13] and the references found there.

The color of an object is a production of the brain and therefore it is not possible to measure color directly. The cause of the color sensation, the electro-magnetic spectrum reaching the eye, is however physically measurable. Processing of color in the human sensory system is done with three different type of photo-receptors, with total range of 380 nm to 780 nm (in practice 400 nm to 700 nm) and sensitivity maxima at 440 nm, 545 nm and 590 nm. The wavelength response curves of these photo-receptors can be found in [14]. Although, the visual system is processing chromato-spatial information we will concentrate here on the color information only. In addition to our interest in the color information as such, our approach is motivated by present models of the human visual system with a division into spatial and chromatic channels.

The spectral nature of the origin of the color information was discovered by Newton and several investigations of the spectral analysis of the color can be found, for example, in [3], [12] and [4]. The three dimensional representation and analysis has dominated color research and applications due to the model of human color vision. These three dimensional representations are mostly based on human tristimulus values

$$X = k\int S(\lambda)\bar{x}(\lambda)\,d\lambda \qquad Y = k\int S(\lambda)\bar{y}(\lambda)\,d\lambda \qquad Z = k\int S(\lambda)\bar{z}(\lambda)\,d\lambda$$

where $S(\lambda)$ is the input spectrum, $\bar{x}(\lambda), \bar{y}(\lambda)$ and $\bar{z}(\lambda)$ are the sensitivity functions of the human eye and k is a normalizing factor. There are thus an infinite number of different spectra with the same values of X, Y and Z.

In the study of color processing it is of great interest to know how many basis vectors (corresponding to the sensitivity functions) are needed for representing the color information and what the optimal basis vectors for three dimensional color representation are. Investigations of the first problem can be found in [6] and [12].

A study of these problems requires access to the whole color information, and we will therefore use the color spectra as initial information. We will use three databases of spectral data. The content of these databases should reflect the properties of the space of those colors that are perceivable by the human visual system. We will first describe the contents of these databases and compare their statistical properties. Then we will use the spectra in these databases as input to a unsupervised neural net and investigate the behavior of the net under different circumstances.

The databases

In our experiments we used three different databases: Two are based on the Munsell Book of Color [9] and one on the NCS system. The old Munsell database consisted of 1253 spectra represented by 61 measurements each. We often ignored the first 5 measurements in these spectra since the illumination in this spectral band was too noisy. The remaining values represent the reflectance in the range from 425 nm to 700 nm with a step length of 5 nm. We refer to this database as the Munsell-I database if necessary. We will call the new database based on the Munsell colors the Munsell database. This database consisted of 1269 spectra of 421 measurements each. They are in the range from 380 nm to 800 nm in 1 nm steps. The NCS database consisted of 1513 spectra based on the NCS color system. The spectral range was 380 nm to 780 nm in 5 nm steps.

In our first experiment we selected from the three databases the measurements in the 425 nm to 700 nm band and computed from them the eigenvectors and eigenvalues of the correlation matrix.

As an example we show in Figure 1(a) the second eigenvector from the three different databases. We found that the first eigenvectors are almost identical for all databases. In Figure 1(b) we see the first three eigenvectors from the Munsell database. Table 1 contains the correlation coefficients between the first four eigenvectors from the Munsell databases and the NCS system. These vectors were calculated from the 425 nm to 700 nm range with 5 nm steps. The eigenvalues from the full Munsell spectra are listed in Table 2.

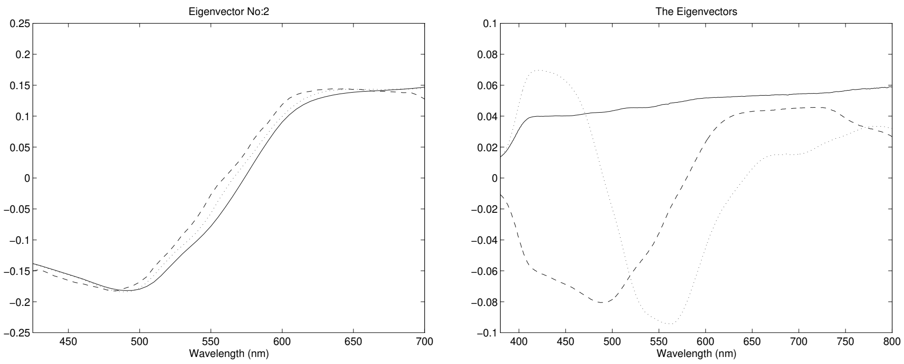

(a) (b)

Figure 1: (a) Second Eigenvectors (b) First three Eigenvectors

Number	NCS vs. Munsell-I	NCS vs. Munsell	Munsell vs. Munsell-I
1	0.9940	0.9982	0.9987
2	0.9847	0.9965	0.9953
3	0.9835	0.9946	0.9942
4	0.9662	0.9939	0.9780

Table 1: Correlations of eigenvectors from the different databases

The conclusion that we can draw from these experiments is that the databases have very similar statistical properties. Therefore we will mainly work with the Munsell databases in the following.

Number	1	2	3	4	5	6	7
Eigenvalue	47.5567	2.3453	0.8268	0.1814	0.0635	0.0331	0.0319

Table 2: Eigenvalues from the full Munsell database

Learning as optimization procedure

Many unsupervised systems in the neural networks literature are specified by their learning rules, i.e. how they change their internal weights. In this paper we will not follow this tradition. Instead we select first some properties of the output vectors computed by the system. Then we design a quality (or energy) function which measures these properties. Finally we derive the learning rules by formulating updating as an optimization procedure. This has the advantage that the huge existing knowledge in optimization theory can be applied to find efficient implementations of the learning procedure. Note that a given energy function can lead to different learning rules depending on which optimization method is used.

In all our investigations we start from the requirement that the extracted feature vectors should convey as much "information" about the space of input signals as possible. The term "information" as used here is not the entropy based definition used information theory but it is used in the Karhunen-Loéve sense of minimum mean squared error. It turns out that a useful class of energy functions can be build around the determinant of the output correlation matrix. We describe this as follows: The patterns in the database are given by the vectors p. The correlation matrix C_p of these signals is $C_p = E(p \cdot p')$ where the prime denotes transposition and $E(x)$ is the expected value of the stochastic variable x. The output of the system for the input p is given by the vector $o = f(p)$ and the correlation matrix of the output is $C_o = E(o \cdot o')$. We choose to select as the first term in our energy function the determinant of the output correlation matrix (see [7] for further information):

$$Q_D = -\det(C_o) \tag{1}$$

In the case where the function f is linear we have $o = W \cdot p$ with a weight matrix W. In this case it can be shown that the matrices W which lead to an extremum of Q_D are connected to the first eigenvectors of C_i via a rotation. This implies that we can multiply W by a rotation matrix and still get the same energy value. In our previous investigations (see [8], [10], [11]) we used this freedom and introduced new requirements for the output vectors. These additional properties were introduced by modifying the energy function to functions of the form $Q = Q_D \cdot Q_A$ where Q_A is an additional term. We found that terms depending on fourth order statistical moments of the output values produced especially interesting results. At about the same time and independent of it, the importance of fourth order moments in neural network research was recognized in the context of sparse coding [5]. The variant that worked best and that was easiest to implement is given by the variance of the squared output values (Q_V) but we experimented also with other variants like Q_4 both described in Equation 2 (note however that the Q_4 term requires that the output values are in the range $[-1, 1]$).

$$Q_A = Q_V = \sum_{n=1}^{N} \mathrm{var}\left(o_n^2\right)) \qquad\qquad \frac{1}{Q_A} = Q_4 = \sum_{n=1}^{N} E\left(o_n^2(1-o_n^2)\right) \tag{2}$$

In many applications it might be necessary to impose certain restrictions on the values of the weights in the system. An example is the optical implementation of the filters where it is desirable to have only positive filter coefficients. Other useful conditions are weight vectors of unit length or weights in the interval $[0,1]$.

These boundary conditions are relatively easy to incorporate into the optimization process. In our implementation we enforce unit-length weight vectors by using polar coordinates. Optimization is then done in the angular variables. In a similar way we enforce positive weights by using weights of the form $w_{ij} = \mathrm{abs}(x_{ij})$. Optimization takes place in the $x-$variables but computation of the output values in the $w-$variables.

Experimental results

In our experiments we used two different optimization programs. In one series of experiments we used an implementation of the CONMIN algorithm together with the line-search algorithm developed by Moré and Thuente. (For more information see the references in [10]). We found that the optimization was very efficient. Using these algorithms has furthermore the advantage that they can be implemented in parallel systems in which the different modules are updated in parallel.

In the first series of experiments we compared the performance of systems based on the energy functions Q_D/Q_4 and $Q_D \cdot Q_V$ (Equation 2). In these experiments we trained a system consisting of three units. We used the CONMIN optimization program and described the weights in polar-coordinates, thus enforcing the strict boundary condition that the rows of the weight matrix should have unit length. We used the Munsell-I database and normed each individual entry in it to length one since the function Q_D/Q_4 requires outputs in the range $[-1,1]$. The result is shown in Figure 2.

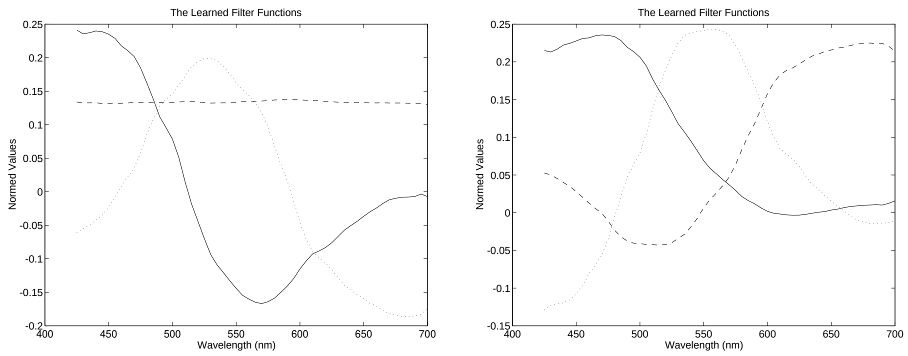

(a) (b)

Figure 2: (a) CONMIN using Q_4 and (b) CONMIN based on Q_V

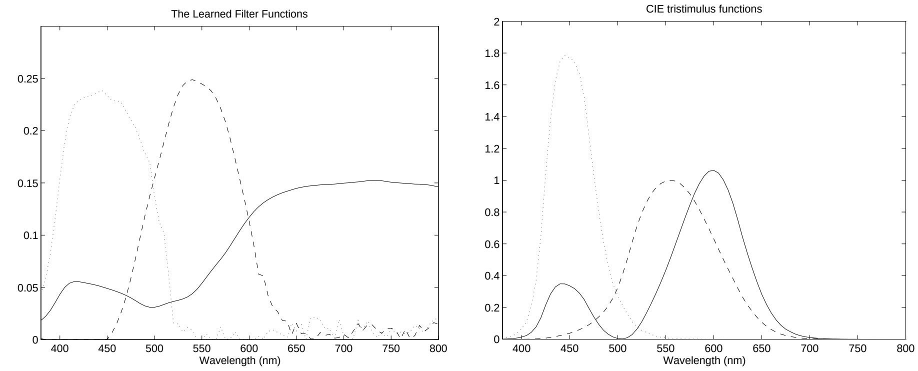

(a) (b)

Figure 3: (a) CONMIN using Q_V and positive weights, (b) Tristimulus functions

These figures show that the system based on the Q_4 function stabilizes in a state that is qualitatively similar to the three eigenvectors of the correlation matrix (Note however that each input vector was normed to length one in these experiments). The system based on the Q_V function on the other hand tries to split the domain into three separate spectral bands.

In another experiment we used the unnormed entries from the new Munsell database and the CONMIN program together with the Q_V function. Here we optimized in the polar coordinates but we used as weights the absolute values of the cartesian coordinates $x : w = \text{abs}(x)$. This enforced positive weights. The resulting filter vectors are shown in Figure 3(a).

This shows that the positivity constraint enhances the ability of the original system to partition the spectrum in different bands. The resulting filters could be characterized as red-green-blue filters. These filter functions can be compared with the tristimulus functions $\bar{x}, \bar{y}$ and $\bar{z}$ mentioned in the introduction. These functions are shown in Figure 3(b) and we see that the main difference lies in their behaviour in the region with the long wavelengths: here the tristimulus functions fall to zero but one of the learned filter functions has relatively high values in this region. An inspection of the spectra in the database shows that the spectra have comparatively high values in this region and since the filter functions are designed to have high correlations with the input patterns it is understandable why they have high coefficients in this region. It is however still an open question why the human visual system chooses to ignore information in this spectral region. We also did some experiments in which we first multiplied the original spectra with the values of the $\bar{y}$ function of the tristimulus system (the dashed curve) shown in Figure 3. These values are similar to the spectral sensitivity of the human visual system. Then we trained our system with this new dataset. The result of one such experiment is shown in Figure 4(a). We see that there are clear similarities between the learned system and the tristimulus system. The main similarity is perhaps that there is one bimodal curve and two bandpass filters. We can however also notice that the learned filters are more concentrated in the middle region which is probably a consequence of the concentration of the $\bar{y}$ curve in this region.

The computation of the output vectors o (or the X, Y, Z-values) mentioned in the beginning is a projection from the high dimensional space of the spectra to a low-dimensional coordinate space. Obviously some information is lost in this process and there are therefore many different spectra that map to the same coordinate vector. In many applications it would however be interesting to compute from the coordinate values (for example the o or the (X, Y, Z) vector) an approximation of the underlying spectrum. This approximation is easy if the filter functions are the eigenvectors since these vectors form an orthonormal system. It is more complicated for the learned filter systems or the X, Y, Z system since they are not orthogonal. It is however also possible for non-orthogonal systems to compute such approximations and in Figure 4(b) we summarize the results of some of our experiments. The figure shows the performance of the approximations based on the eigenvector system (o), the (unrestricted) filter systems learned from the Munsell database ($\times$) and the learned positive filter functions ($+$) learned from the same dataset. The error was computed as the mean value over the quantity: $100 \times$ (reconstruction error/signal energy). We see that the learned filter systems and the eigenvector systems are comparable whereas the system with the positivity constraint has a slightly worse performance.

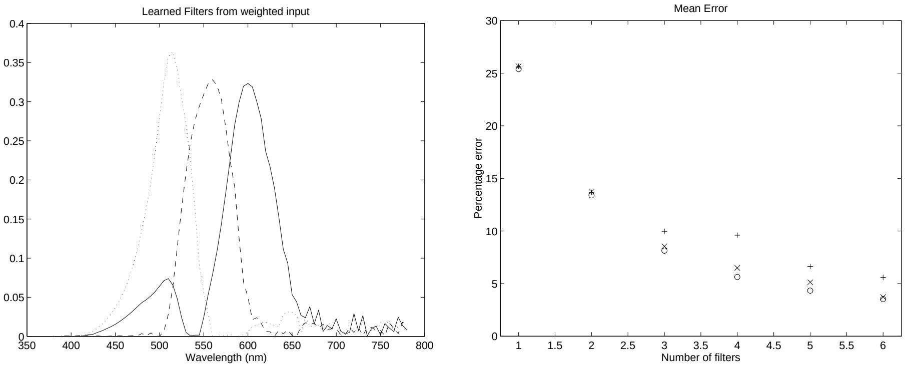

(a) (b)

Figure 4: (a) Learned filters for weighted input (b)Approximation errors

Conclusions

We investigated three databases containing reflectance spectra describing the colors relevant for human perception. We then introduced unsupervised neural networks whose learning rule was derived from energy functions that included second and fourth order statistics of the output vectors and trained them with these reflectance spectra. We also demonstrated how additional constraints (like the positivity of the weights) can be incorporated into the network design.

We found that the performance of the learned filter systems was in general comparable with the eigenvector system. This is also true for the systems in which we required that all filter coefficients should be non-negative. These filters are especially interesting for optical realizations of filters. Due to it's parallel nature, optical computing is an interesting option for operations in the optical part of the spectrum.

Acknowledgment

Reiner Lenz and Mats Österberg were supported by the TFR-grant TFR-93-192. The NCS data base was obtained from the Scandinavian Color Institute in Stockholm and provided by B. Kruse. J. Hiltunen is supported by the Machine Vision program of TEKES. J. Hiltunen and M. Österberg received travel grants from NORVIC.

References

[1] Atick, J. J. Could information theory provide an ecological theory of sensory processing? *Network 3*, 2 (May 1992), 213–251.

[2] Attneave, F. Some informational aspects of visual perception. *Psychological Reviews 61* (1954), 183–193.

[3] Cohen, J. Dependency of the spectral reflectance curves of the Munsell color chips. *Psychon Science 1* (1964), 369–370.

[4] D'Zmura, M., and Iverson, G. Color constancy. iii. general linear recovery of spectral descriptions for lights and surfaces. *JOSAA 11*, 9 (1994)

[5] Field, D. J. What is the goal of sensory coding? *Neural Comp.*, 1994.

[6] Jaaskelainen, T., Parkkinen, J., and Toyooka, S. Vector subspace model for color representation. *JOSAA 7*, 4 (April 1990), 725–730.

[7] Lenz, R., and Österberg, M. Computing the Karhunen-Loeve expansion with a parallel, unsupervised filter system. *Neural Comput./* 1992.

[8] Lenz, R., and Österberg, M. A new method for unsupervised linear feature extraction using fourth order moments. *Pat. Rec. Letters*, 1992.

[9] Munsell Color. *Munsell Book of Color, Matte Finish Collection*, 1976.

[10] Österberg, M. *Quality Functions for Parallel Selective Principal Component Analysis.* PhD thesis, Linköping University, 1994.

[11] Österberg, M., and Lenz, R. Unsupervised parallel feature extraction from first principles. In *Neural Information Processing Systems 6, 1994*.

[12] Parkkinen, J., Hallikainen, J., and Jaaskelainen, T. Characteristic spectra of Munsell colors. *JOSAA 6*, 2 (1989), 318–322.

[13] Taylor, J. G., and Plumbley, M. D. Information theory and neural networks. In *Math. App. Neural Networks*, J. G. Taylor, Ed., 1993.

[14] Wyszecki, G., and Stiles, W. *Color Science*, 2 ed. Wiley & Sons, 1982.