Decomposition of a Mixture of Signals in a Model of the Olfactory Bulb*

O. HENDIN and D. HORN
School of Physics and Astronomy Raymond and Beverly Sackler Faculty of Exact Sciences Tel Aviv University, Tel Aviv 69978, Israel
J. J. HOPFIELD
Divisions of Chemistry and Biology California Institute of Technology Pasadena, California 91125

ABSTRACT We describe models for the olfactory bulb which perform separation and decomposition of mixed odor inputs from different sources. The odors are unknown to the system, hence this is an analog and extension of the engineering problem of blind separation of signals. The separation process makes use of the different temporal fluctuations of the input odors which occur under natural conditions. We discuss two possibilities, one relying on a specific architecture connecting modules with the same sensory inputs, and the other assuming that the modules (e.g. glomeruli) have different receptive fields in odor space. We compare the implications of these models for the testing of mixed odors from a single source.

INTRODUCTION

The olfactory system is most suitable for the application and development of neural network models. This system is known to be one of the oldest phylogenetic neural systems in mammals. Its architecture seems to be relatively simple, yet its detailed functional behaviour is still an open and challenging problem [1].

Candidate membrane receptors for odor molecules were recently identified [2]. They activate the olfactory sensory neurons which, in all species, project directly to glomeruli. The latter [3] compress the incoming information and form the first encoding stage in the olfactory pathway. The glomeruli, together with excitatory mitral/tufted cells, and inhibitory interneurons, belong to the olfactory bulb. The processing of information within the bulb is taking place mainly via dendro-dendritic interactions. The output is carried by mitral/tufted neurons to central regions of the brain. We will be concerned with the first computational stages of the bulb, performing an analysis of the incoming information. We expect the bulb to be responsible for separation of mixed odors, classification in terms of their major components and, possibly, identification of odors by comparing them with memories in the olfactory cortex [9] [10]. The latter function, which has to rely on feedback connections, will not be discussed in this paper.

Separation of mixed odors can be achieved by computations which make use of independent temporal fluctuations in the odors which comprise the mixture. This technique is known as 'blind separation' in the literature of signal processing [5] [6] and can be implemented in a neural network with adaptive linear connections [4]. One starts with a network of mitral cells, each receiving some input from the sensory neurons, thus spanning an effective odor space. The adaptive learning takes place with each new environment, leaving after some time only two active neurons when, for example, two independent odors exist in the environment. These neurons correspond to the leading components of each of the two odors. The interesting feature of this model is that it succeeds in performing the separation without having to invoke any prior knowledge of the odors which are being analyzed. Nonetheless, it is not yet suited for classification purposes because it activates only the neurons which correspond to the leading components of each of the mixed odors.

Our interest lies in generalizing this model in such a fashion that the separation process will take place but will also transmit information about different components of the input odors. We will present two models which lead to such results. Both involve multiple units of the kind described above, to which we refer as replicas. The first model assumes that all replicas have the same input, but have a dynamical hierarchy between them. We will refer to this model as "vertical replicas" implying the hierarchy. The second model has no connections between the replicas but each one receives a different input, i.e. views a different section of odor-space. We refer to it therefore as a model of horizontal replicas.

1 The Basic Module

The olfactory sensory neurons converge onto the glomeruli and terminate on the dendrites of mitral/tufted cells. We regard the latter as representing odor space. To be specific, let us choose a set of mitral cells labelled by $i = 1, \cdots, N$ to represent this vector-space. Since it is known that the excitatory mitral cells interact in a dendro-dendritic fashion with inhibitory granule cells, we follow Ref. [4] and describe this interaction by an effective matrix T_{ij} in a linear differential equation for the mitral cell potentials u_i:

$$du_i/dt = -u_i/\tau - \sum_j T_{ij} u_j + I_i(t) \qquad (1)$$

I_i represents the input which reaches cell number i. This input is composed of signals from $j = 1, \cdots, K$ different sources which are assumed to have independent temporal fluctuations $a_j(t)$:

$$I_i(t) = \sum_{j=1}^{K} S_i^j a_j(t) \qquad (2)$$

The parameters S_i^j represent the composition of the odor number j in the vector space spanned by the

*Acknowledgment Work at Caltech was supported in part by the Ronald and Maxine Linde Venture Fund.

mitral cells. The statistical independence of the odors $a_j(t)$ allows for performing blind separation. This is brought about by a learning rule for the effective inhibitory synapses:

$$dT_{ij}/dt = \bar{u}_i \bar{u}_j [\delta + \epsilon(\bar{u}_i - \gamma \bar{u}_j)] \quad (3)$$

The bar notation refers to high-pass filtering of the potentials. The procedure which we employ removes the average value of u_i and leads to $-1/2 < \bar{u}_i < 1/2$ which has all high frequency fluctuations of the original u_i. The matrix element T_{ij} changes with time according to 3 but is not allowed to be negative, i.e. negative contributions to 3 are ignored if T_{ij} reaches zero.

For appropriate choices of the parameters which appear in these equations, the system can separate the different inputs. In the case $N = 6$ $K = 2$ Hopfield [4] has demonstrated how two of the six neurons become dominant, each fluctuating in time as one of the input odors. This is also shown in Fig. 1. The last term in 3 is responsible for just one component surviving for each odor. This will be the largest component in the vector representation of the odor. We have here a winner-take-all algorithm within each odor vector.

2 Model of Vertical Replicas

Let us consider a system of L different replicas of the basic module which inhibit one another:

$$du_i^n/dt = -u_i^n/\tau - \sum_j T_{ij}^n u_j^n - \lambda \sum_{m=1}^{n-1} u_i^m + I_i(t) \quad (4)$$

$n = 1, \cdots, L$ The replica number is given by the upper index. Inhibition occurs in a hierarchical manner, where each replica is inhibited by all previous ones. As a result, replica number 1 will be identical with the model described by equations 1 and 3. The others obey different dynamics, influenced by previous replicas. T_{ij}^n continue to obey 3 , however the cell potentials which drive the learning procedure are the ones corresponding to replica number n .

This construction yields a network of size $L \times N$ in which every sub-net (replica) contains N neurons interconnected via dynamic connections T_{ij}^n. In addition, each neuron is inhibited by all the neurons at the same location in the previous sub-nets and inhibits its counterparts in the following replicas. This architecture can be viewed as a series of neuronal surfaces, each projecting onto all subsequent surfaces, while all receive the same external input.

Since the adaptive learning rules have the character of a winner-takes-all algorithm within each odor, the first sub-net is expected to activate the leading components in the odors. The inhibition between the sub-nets causes the next one to activate the next-to-leading components, and so on. Clearly ambiguities may arise if two components of an odor are of the same magnitude, or if two odors have the same biggest component. Barring such coincidences, the network performs as expected.

We solve numerically a model with $L = 3$ replicas of $N = 6$ components in odor-space. A superposition of $K = 2$ odors is used as input. After the system stabilizes we find in each replica just two elements which have large amplitudes, thus establishing blind separation of two odors. Fig. 1 shows the amplitudes of the six neurons of replica 1. Neurons number 1 and 6 vary with large amplitudes, whose fluctuations follow those of the two odors used as inputs. Neurons numbers 2,3,4 and 5 have negligible amplitudes. The basic time scale of the system is 1ms. The input signals have a Fourier spectrum ranging between 0-20 Hz with a strong peak at about 6 Hz.

Figure 2 displays the results for replica number 2. For higher replica numbers one has to wait longer until the separation process sets in, since each replica is influenced by all previous ones. A comparison between the leading components of all replicas is shown in Figure 3. The values of the two odor vectors used in these examples were $S^1 = (.99, .72, .62, .44, .32, .23)$ $S^2 = (.29, .37, .47, .60, .77, .99)$. As is quite evident from these figures, the algorithm performs its task leading to the correct conclusion that one odor has components 1-2-3 (where the notation implies decreasing strength), while the other has components 6-5-4.

Several conditions have to be met for the separation process to succeed. One is that the major components of the same odor vector must differ significantly from one another so that the system can establish their order. The components S_i^j of the input vectors in our examples are chosen to lie between 0 and 1. If the difference between major components shrinks to about 0.1, convergence is slowed down to about 10s. This increases to 45s when the difference decreases to 0.05, which is the limit below which separation will no longer work. A second condition is that the differences between the values of the same component in the two different odors should be larger than 0.1. A third condition is that the component whose turn it is to win in a given replica should be bigger than the corresponding component in the other odor. This condition is not absolutely necessary but, if it is not met, success is not guaranteed.

A fourth, and obvious, problem is that if the same component has the same significance in both odors (e.g. is the highest one), separation is impossible. While this is a problem of principle, any other overlap (such as the same component being leading in one odor and second in the other) is just a nuisance. In the example analyzed above, there was no overlap between the major components of the two odors. We have however tested cases with overlap of major components such as 1-2-3 and 2-6-5, or even 1-2-3 and 2-6-1. In both cases our model performs the separation but the convergence times are longer, and increase as more pairs of the same components have to be activated in the replicas.

3 Scaling and Binding

Trying to build a biologically plausible model on the structure described above we now address two important challenges. The first concerns the size of the system: can we scale up the structure to involve more neurons and more odors? (This would give a useful function to size per se of the early processing system.)

Clearly if we were to scale up the network and assume that the odor representation is randomly distributed among all neurons we run into the problem that sooner or later there will be major components whose weights differ by less than 0.05. This may be resolved in the system of horizontal replicas discussed below, but cannot be handled by the vertical system. Moreover, it seems that the resolution time required of the vertical replicas system limits us to an odor description which is characterized by just a few leading components. Therefore, we will consider a larger system on which an odor is represented on the average by four components which lie between 0.5 and 1 and many (insignificant) components whose weights are below 0.5. Under this assumption we obtained blind separation of three odors in a vertical system of 24 neurons. The convergence times were of the order of 7s, 12s, 24s and 34s for replicas numbers 1 to 4. The number of mixed odors which can be separated by the system will grow linearly with the number of neurons. Nonetheless, if the weights are chosen randomly, a mismatch may occur if the conditions specified in the previous chapter are not met.

In the different figures we have identified the odors by comparing the fluctuations of the output with those of the input. The biological system has to do something else. It can, for instance, correlate all outputs which fluctuate in the same way. Moreover, it should be able to do it on the level of the action potentials of the cells whose membrane potentials were the subject of our analysis so far. This is a typical binding problem [11]. It was already demonstrated that binding can be achieved under such conditions by Horn Sagi and Usher [12]: If the odor components were represented by oscillating cell assemblies, the fluctuations could drive them into phase-locking, representing binding [13]. The implementation of our scheme in an oscillatory environment needs further investigation. In principle one can follow two directions: either generalize the present system to include competitive oscillations or allow for another computational step for which our series of outputs serves as the input.

4 Model of Horizontal Replicas

Here we start with the assumption that each replica describes an ensemble of mitral and tufted cells which receive sensory inputs which differ from those of other replicas. We do not require any connections between different replicas for the calculational task at hand. To be specific, we assume that odor space is defined by the projections from the sensory neurons and each replica receives a different section of this space, i.e. has its own "receptive field". We represent this receptive field by a Gaussian with mean location μ and width σ in a space of N_o ordered inputs (odor-space). We allow for the possibility that the Gaussian distributions of different replicas have large overlaps with one another.

Under the influence of a single odor each replica will go through its competitive training procedure and select the strongest component within its own receptive field. Thus an odor will be represented by a set of components selected by different replicas, a representation which is correlated with the input but not identical with it. Examples are shown in Figure 4, where we choose $N_o = 26$ mitral cells and $L = 10$ replicas with μ varying by three steps from one replica to another. The value of σ determines the effective number of sensory cells to which a given replica is tuned. The two parallel lines in this figure represent the locations of $\mu \pm \sigma$ in the different replicas. The inputs into the system are shown in the upper frame, and the representations of the two odors are identified with the replica-mitral matrix of figures 4a and 4b. Note that within this scheme we do not have to limit ourselves to the scaling assumption of the previous chapter. Any number of leading components is allowed, as long as not too many fall in any one replica.

Now the question is whether in the presence of two odors we will be able to reconstruct each of them simultaneously. The answer to that is given in Figure 4c. All but one replica converged to the sum of the representations of the two odors. We have repeated these calculations for various width parameters in order to get a qualitative feeling for the performance of this model. There are two relevant parameters in our calculations. One is the width σ of the Gaussian projections on the replicas. The other is the sparsity s of the input vectors. s represents the probability that an odor component will be different from zero (whereupon it is distributed randomly between 0 and 1). Clearly small s guarantees low overlap between different odors and, therefore, better reconstruction of their superposition. This is further improved by choosing also low σ. These expectations are borne out by the results shown in Table 1, where the number of misidentifications are shown. We find that, over a wide range of parameters, the mitral cells which are being activated by the model for the mixtures of two odors are the same as those activated by the pure odors within an accuracy of 10% – 20%.

$s\sigma$	3	5	7
0.7	1.2 ±1.2	1.8 ± 1.1	1.7 ±1.0
0.6	1.4 ±1.5	1.7 ± 1.1	2.2 ±1.6
0.5	1.2 ±0.8	1.2 ± 1.2	1.4 ±1.2
0.4	1.1 ±0.1	0.6 ± 0.7	1.4 ±1.3

Table 1

Errors of misidentification in the replica-mitral representation of mixed odors in 10 trials

5 Biological Implementation

The model of vertical replicas solves the problem, keeping track of and sending out signals which describe odor quality information in the presence of multiple odor sources. There remain, however, many open questions concerning whether or how such a scheme is implemented in neurobiology. First, is it plausible that the olfactory bulb indeed is laid out in the fashion required by replica networks? If the different replicas are to be thought of as different sets of mitral cells, it would be necessary to carry specific inhibitory information between corresponding mitral cells (i.e., mitral cells which have the same fundamental response to individual odor components) in adjacent parts of the olfactory bulb. This requires high specificity of connectivity between classes of cells, and thus may seem unlikely. Yet such a specificity would not be much different in character from that which is necessary in the visual system to ensure that neurons carrying color information from different spatial locations can be consistently brought together to evaluate the color of an object.

An alternate possibility is that the piriform (olfactory) cortex performs the computation, for the architecture of this area has inputs from the olfactory bulb extended across a sheet of cortex, and each band of pyramidal cells laid out in a strip perpendicular to the incoming axon collaterals could be thought of as a replica.

It is much simpler to suggest a biological implementation of the model of horizontal replicas. The layer of glomeruli seems to be a natural candidate for such a structure. They transfer the sensory inputs to groups of tufted and mitral cells and seem to differ from one another by the input they receive [3]. In this sense each glomerulus can be viewed as a filter defined on odor-space, in a similar fashion to our theoretical structure. Such an architecture would be in accord with the topographicallly distinct patterns expressing odorant receptor genes within the olfactory epithelium, assuming they are preserved in the axonal projections to the olfactory bulb [14].

Both models require synaptic change over the time scale of several seconds. Experiments searching on this time scale for synaptic change in the glomeruli, in the olfactory bulb, and in the piriform cortex are particularly important to the question of whether theories of this general type might be relevant to early olfactory processing, regardless of the details of particular implementation or information coding.

6 Mixed Odor Tests

The two models which we have discussed have different representations of odors. The vertical model represents an odor by an array of ordered strength of components, whereas the horizontal model represents an odor by the replica-mitral matrix of the type shown in Fig. 4. Moreover, the long time of convergence of the vertical model seems to imply that odor identification is restricted to its few leading components, while in the horizontal model there is no size limit.

These differences have interesting consequences. Consider what happens if the two leading components of an input odor switch their relative strength. In the vertical model this implies changing the identity of the odor, whereas in the horizontal model it could be regarded as a small change which still leaves most of the representation intact. The vertical model thus would have built-in category boundaries at an early olfactory processing stage.

To show an example we choose an input based on a mixture of two odors with the same temporal structure:

$$I_i(t) = (xS_i^A + (1-x)S_i^B)a(t) \qquad (5)$$

where $0 < x < 1$. The question of interest is how will the system switch from the identification of odor A to that of odor B as the mixing parameter x is varied. In Fig. 5 we show the results of the horizontal model for some random choice of odors. We plot the fraction of matrix components that agree with those of the pure odors. As expected, we find a smooth decrease as a function of the degree of mixing.

The example discussed here is quite relevant in the context of psychophysical experiments. It has been demonstrated that humans have a very poor ability to identify odors in mixtures [7] [8] [15]. The mixtures considered in these experiments were prepared through an olfactometer with several channels, hence they correspond to a problem of a single source in our nomenclature, just of the type discussed above.

Many models of perception exhibit sharp categorical boundaries at the perceptual level. Within this context, such boundaries in odor perception are quite different within the two models. In the vertical model, categorical effects will be seen in the collective electrophysiological response of cells in the olfactory bulb (or perhaps the piriform cortex). The overall state of the bulb reflects a categorical decision. In the horizontal model, each mitral cell will exhibit a categorical boundary, but the boundary will be different for each cell, and psychological categorization (presuming it exists) must be constructed higher in the olfactory system. The predicted correspondence (or lack of it) of categories seen in the response of individual cells with the categories of perception is different for the two models.

While psychophysics alone cannot distinguish between these two different implementations of fluctuation-based odor separation, it could be pivotal to understanding whether or not such odor separation is done. Most olfactory psychophysical experiments have been done with complete mixing, in which only one "odor object" is present at a time. Odor discrimination experiments in which different odors come from two or more sources, in a way which guarantees variations in their amplitudes at a subject, would better correspond to natural environments. They are fundamental to the identification of the computation done by the early olfactory system.

FIGURE CAPTIONS

1. The numerical solution of eqs. and as a function of time. The units are in msec. The parameters were chosen as $\delta = 0.005$ $\epsilon = 0.022$ $\gamma = 1$. Shown are the different components of the odor vector space, each represented by temporal variation of one mitral cell. The inputs $a_j(t)$ of the two odors are shown in the top windows, dominated by fluctuations of order 6 Hz. Separation sets in after a few tenths of a second.

2. The amplitude variations in replica 2 whose dynamics is given by with $\lambda = 0.65$. Note the clear inhibition of the leading components of replica 1 (numbers 1 and 6). Separation sets in at about 10s.

3. Comparison of the leading components of all vertical replicas. The notation n-m implies replica number n and neuron number m.

4. Replica-mitral matrix representation of odors in the horizontal replica scheme. Two odors with 26 components are represented schematically in the top frame where the size of the circles are proportional to the strength of the component. These components have random finite values with probability $s = 0.7$. They are projected via Gaussian filters onto the replicas whose dynamics follow the competitive mechanism of section 2 and converge onto the largest filtered component in every replica. The results for the two odors separately are represented in frames a and b. The lines are drawn through the values of $\mu \pm \sigma$ of the different Gaussian filters, representing the receptive fields of the horizontal replicas. Frame c represents the result of the horizontal replica system when these two odors are processed simultaneously. The two cases with half-black and half-white circles are ones in which the largest components of each odor hit the same mitral cell. Separation was therefore not achieved, and this cell remained the dominant one, and continued to exhibit mixed behavior.

5. Test of identification of the system (6.1) by the horizontal model. Odors A and B were chosen randomly with $s = 0.7$, thus guaranteeing some overlap between them. Plotted is the fraction of matrix components that agree with those of the pure odors A or B respectively.

REFERENCES

1. J.L.Davis and H.Eichenbaum (eds),"Olfaction: a model system for computational neuroscience" MIT press (1991).
2. L. Buck and R. Axel, Cell 65, 175-187 (1991)
3. G. M. Shepherd, Nature 358, 457-458 (1992)
4. J. J. Hopfield, Proc. Natl. Acad. Sci. USA 88, 6462-6466 (1991).
5. J. Herault and C. Jutten, AIP Conf. Proc. 151, 206-211,(1986)
6. C. Jutten and J. Herault, Sig. Proc. 24, 1-10 (1991).
7. D. G. Laing, P. Panhauber, M. E. Willcox and E. A. Pittman, Physiol. Behav. 33, 309-319 (1983).
8. D. G. Laing and G. W. Francis, Physiol. Behav. 46, 809-814 (1989).
9. L.B.Haberly, Chemical Senses 10 219–238 (1985).
10. E.D.Kanter and L.B.Haberly, Brain Research 525 175–179 (1990).
11. A.M. Treisman & G. Gelade, Cog. Psychol. 12, 97–136 (1980).
12. D. Horn, D. Sagi and M. Usher, Neural Computation 3, 509–524 (1992).
13. C. von der Malsburg & W. Schneider, Biol. Cybern, 54, 29–40 (1986).
14. K. J. Ressler, S. L. Sullivan and L. B. Buck, Cell 73, 597–609 (1993).
15. M. Laska and R. Hudson, Chem. Senses 17 403–415 (1992).

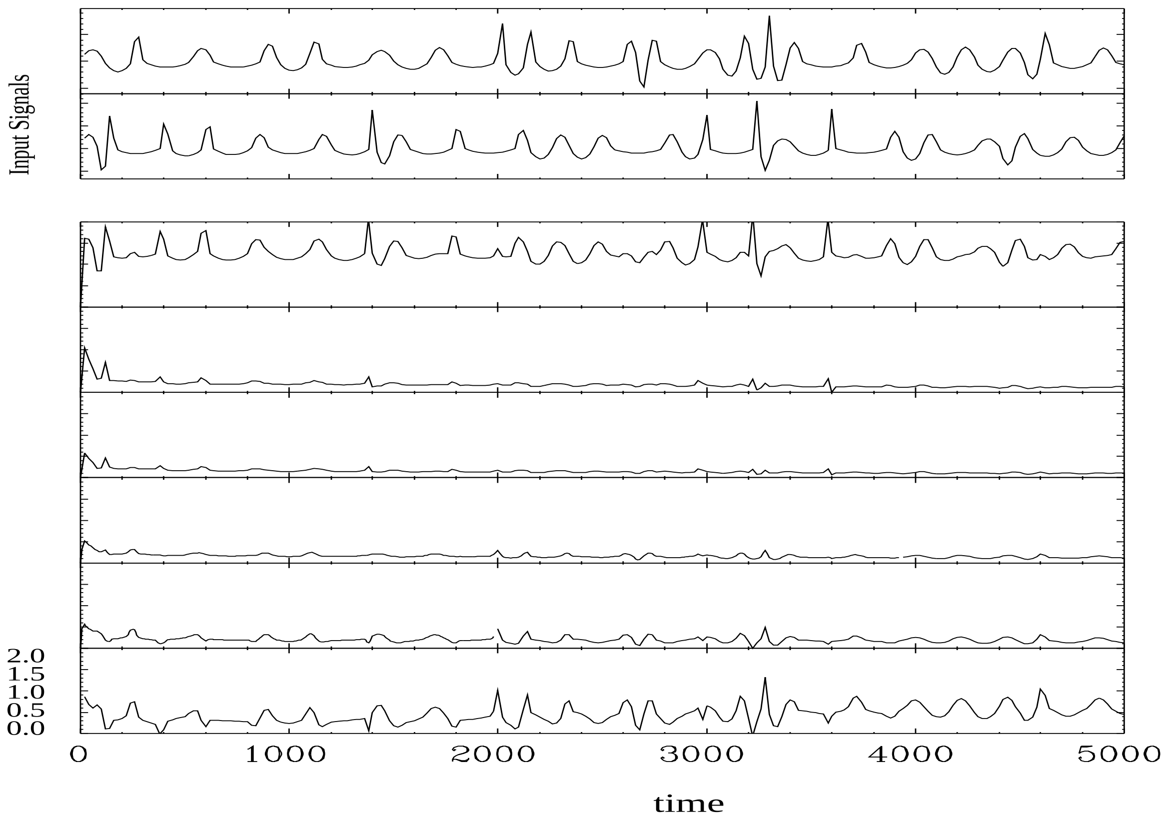
Input Signals
N=1
2
3
4
5
6
2.0
1.5
1.0
0.5
0.0
0
1000
2000
3000
4000
5000
time

Fig. 1

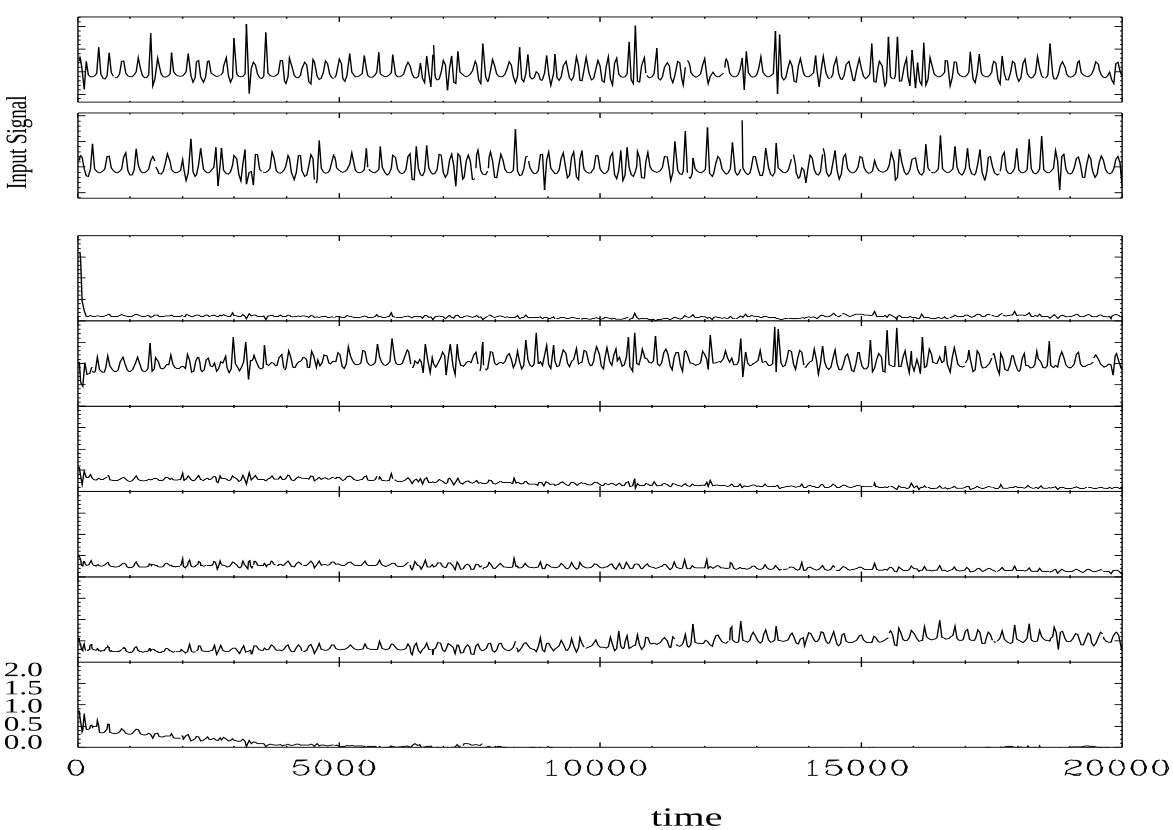
Input Signal
N=1
2
3
4
5
6
2.0
1.5
1.0
0.5
0.0
0
5000
10000
15000
20000
time

Fig. 2

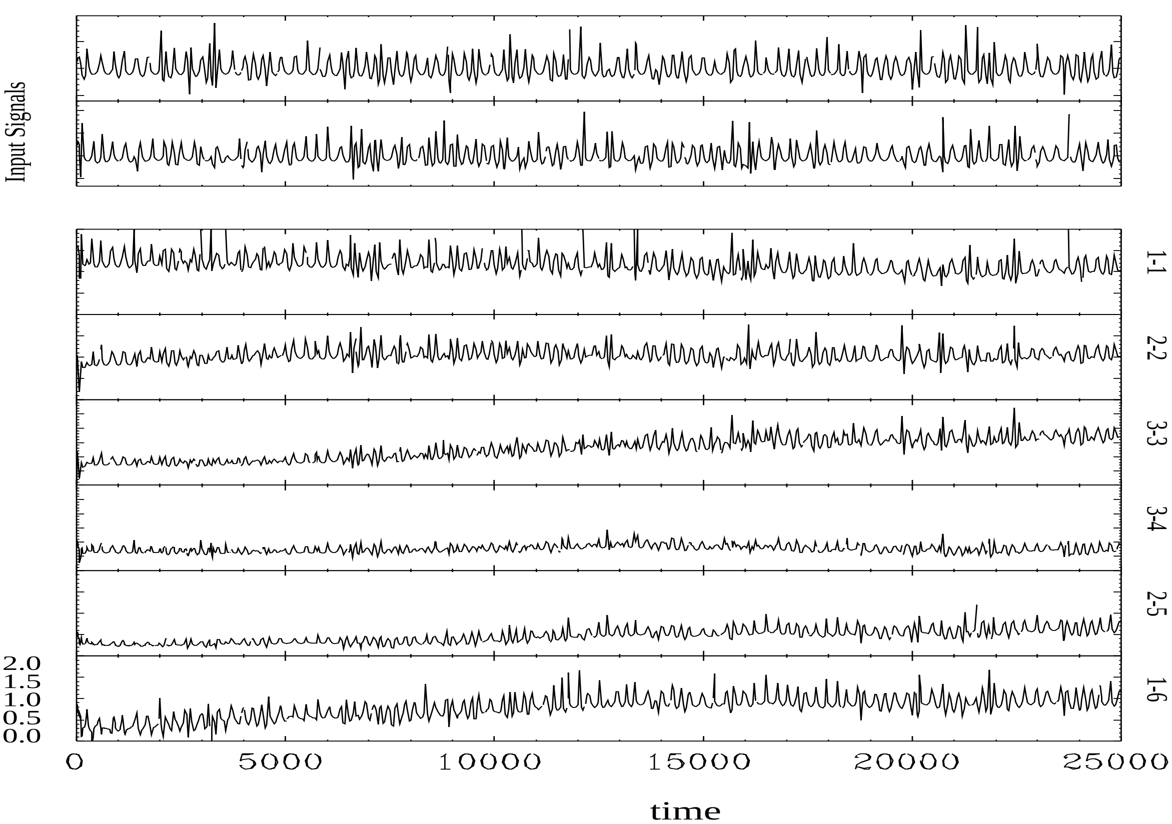
Input Signals
1-1
2-2
3-3
3-4
2-5
1-6
2.0
1.5
1.0
0.5
0.0
0
5000
10000
15000
20000
25000
time

Fig. 3

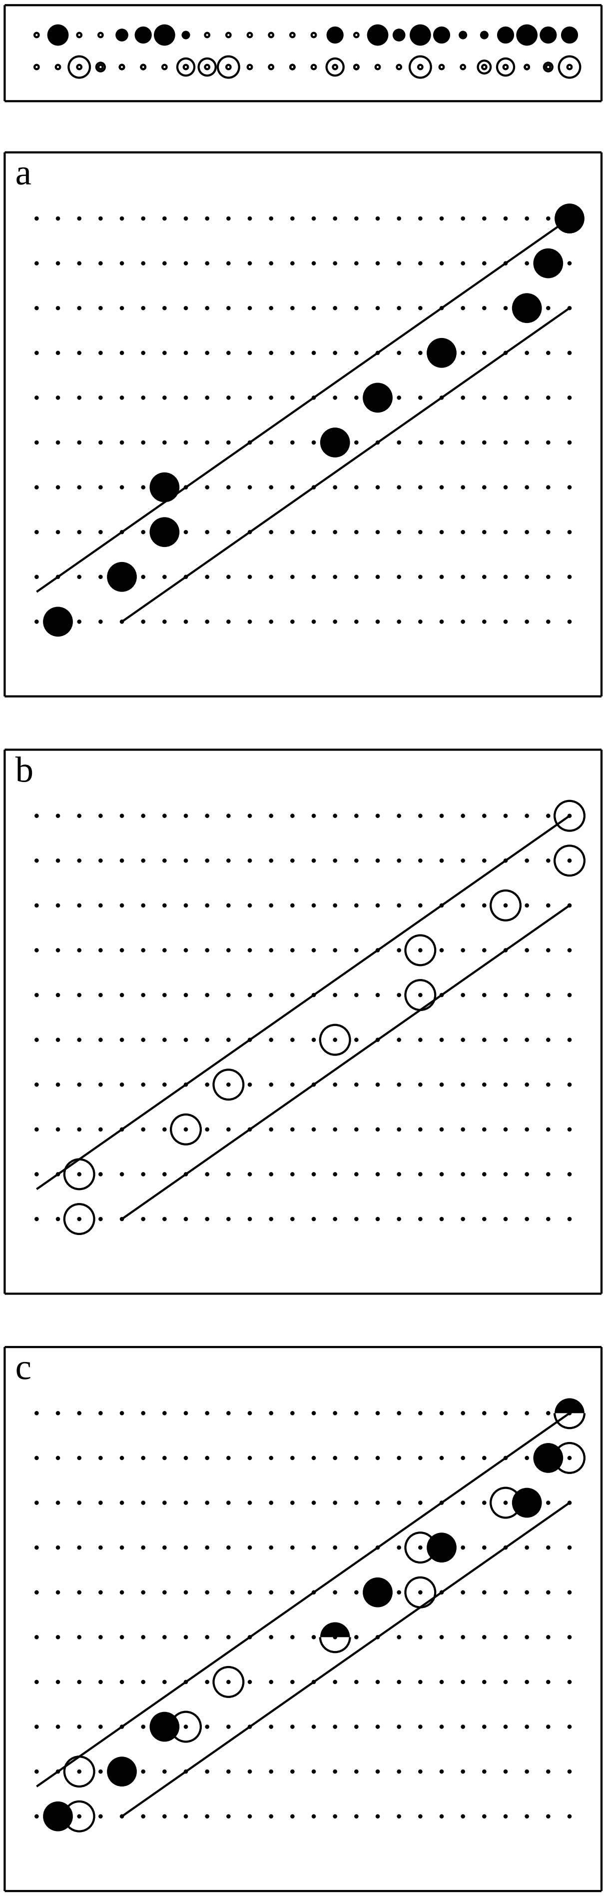

Fig. 4

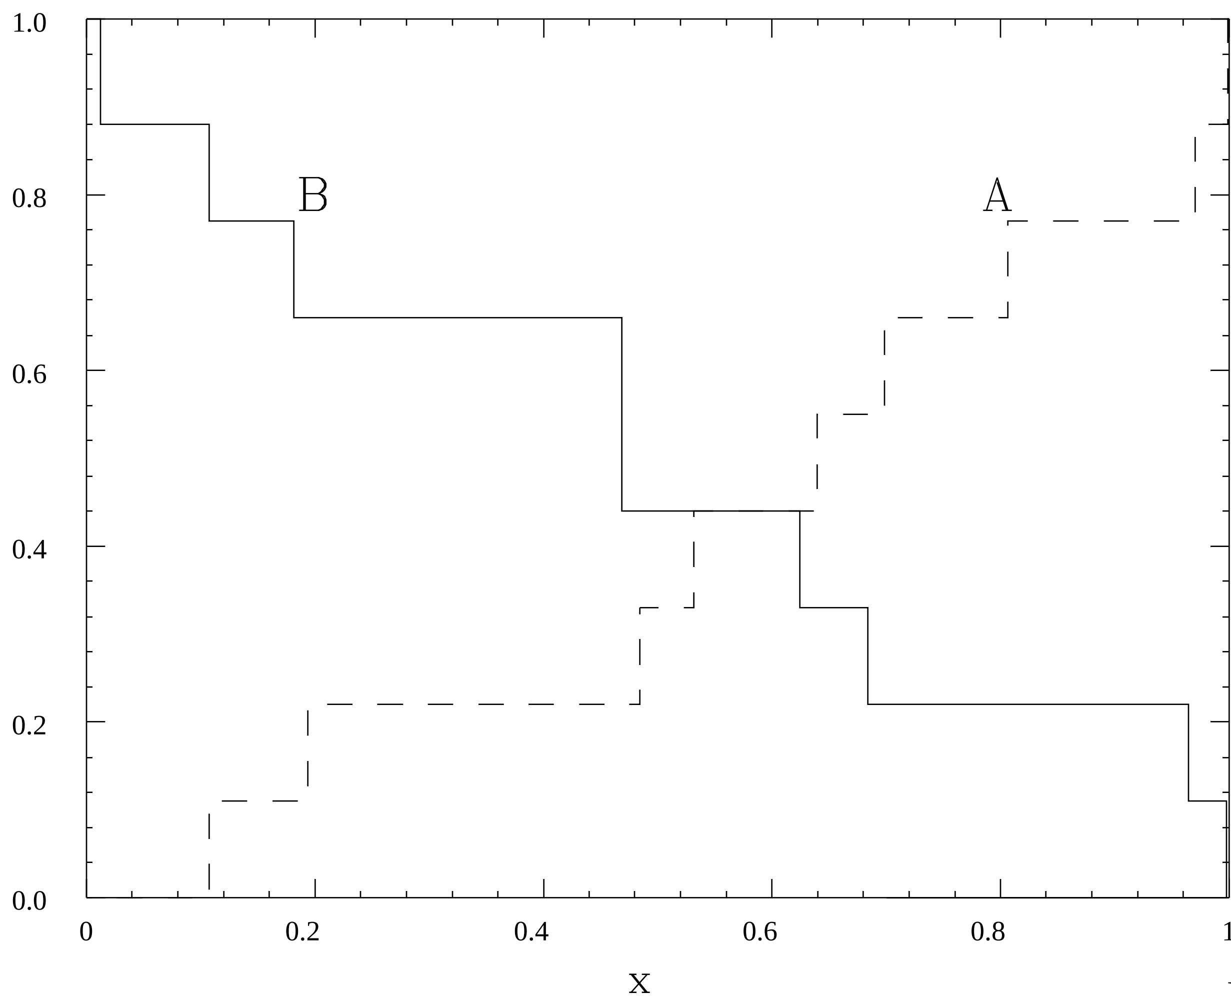

Fig. 5