

Det här verket har digitaliserats vid Göteborgs universitetsbibliotek.
Alla tryckta texter är OCR-tolkade till maskinläsbar text. Det betyder att du kan söka och kopiera texten från dokumentet. Vissa äldre dokument med dåligt tryck kan vara svåra att OCR-tolka korrekt vilket medför att den OCR-tolkade texten kan innehålla fel och därför bör man visuellt jämföra med verkets bilder för att avgöra vad som är riktigt.

This work has been digitised at Gothenburg University Library.
All printed texts have been OCR-processed and converted to machine readable text.
This means that you can search and copy text from the document. Some early printed books are hard to OCR-process correctly and the text may contain errors, so one should always visually compare it with the images to determine what is correct.






Gothenburg Psychological Reports

University of Gothenburg
SWEDEN

Helge Malmgren

Learning by Natural Resonance

Number 6 Volume 21 1991



Learning by Natural Resonance

Helge Malmgren
Department of Philosophy
University of Göteborg

Malmgren, H. Learning by natural resonance. *Göteborg Psychological Reports*, 1991, 21, No. 6. A new model for learning in natural and artificial neural networks is proposed. The system described here can learn associations to sequences of inputs without ad-hoc assumptions about temporal coding. The model is related to adaptive resonance theory and to the theory of adaptive neural oscillators with backpropagation but its teaching principle and its concept of matching between input and output signals are simpler and biologically more plausible. The basic learning concept is that of natural resonance, i.e. the auto-selection of stable or pseudo-stable limit cycles given an alternation between an input constraint and relaxation periods. Such natural resonance can be shown to occur in systems with feedback resulting from a certain kind of non-linear interaction between input and output even if one does not assume any external teacher or an explicit reset-on-mismatch mechanism. Some discrete system simulations which support the theory are reported.

Keywords: Learning, Networks, Adaptive resonance, Automata, Robotics, Neural oscillators.

Introduction and background

MacKay (1956) outlined a model for automaton learning which involves the comparison between the environment and an internally produced representation. An error signal is sent as input if a mismatch occurs. If the environment is reasonably stable this error signal (or repeated such signals) may lead to an *adaptation* in the sense that the internally produced representation eventually matches the environment. In cognitive psychology, this general scheme has become known as the "analysis-by-synthesis" model. As MacKay himself points out, the model can be developed in several directions depending (among other things) on how one conceptualizes the nature of the error signal.

Among recent theories about learning in neural networks which conform to MacKay's conception, Grossberg's "adaptive resonance" theories (ART; cf Grossberg 1987) should be mentioned. The ART models are systems built of continuous, non-linear, neuron-like elements with modifiable synapses. In a matching process, the output ("top-down signal") from the higher-level neurons is compared to the input ("bottom-up signal") and if there is a mismatch an error ("reset") signal is sent. Grossberg's models are very detailed and also include a lot of

This work was supported in part by the Swedish National Board for Technical Development. I am grateful to several colleagues, especially Björn Haglund, Holger Wigström and Olof Östensson, for helpful criticism.

other features as for example the *vigilance* parameter which determines the sensitivity of the reset mechanism to mismatch.

Like most recent neural network theories, the ART models are built around the concept of a system with stable attractor states. This means, among other things, that learned associations are stored as stable parameters of the connections between neuronal elements. Although there are many useful mathematical results concerning stability and convergence in such systems (mainly related to the concept of an energy function) it can be argued that they cannot explain on-line learning of dynamical patterns in a natural way. The theory of adaptive neural oscillators therefore takes its departure in the more general concept of stable (or approximatively stable) limit cycles. In a recent version of these models (Doya & Yoshizawa 1989), continuous back-propagation of error is used as the basic teaching principle. A multi-layered system of non-linear neuronal elements is first exposed to a dynamical input having a certain waveform. The mismatch between this input and the system's output determines the nature of the error signal which is continuously fed back to the system according to a version of the back-propagation algorithm and which tends to correct the output signal in the desired direction. In a relaxation period, the system then takes its own output as input. After repeated alternation of training and relaxation the system often reproduces the original waveform very closely. The main drawback of the theory is - as the authors themselves remark - that the back-propagation algorithm is extremely implausible as a biological principle.

In this article, an alternative approach to the learning of dynamical input is presented. Like the theory of adaptive neuronal oscillators it relies on the concept of stable or pseudo-stable limit cycles, but unlike that theory the present model uses unsupervised learning, i.e., it does not involve any assumptions about external teaching mechanisms. In the latter respect the present model is more similar to adaptive resonance theory but it differs from ART in that it dispenses with the assumption of stable attractor states and in that the error signal and the matching of top-down and bottom-up signals are modeled in a radically different way.

The model is intended to be applicable to all discrete systems and to a large set of continuous systems. In this article we will illustrate our theory by means of a simple class of applications, namely that of *randomly composed finite automata with input*. These devices are completely deterministic but have a transition matrix which has been constructed in a random way, except possibly for a certain bias towards stability, in which case we call them *inert random automata*. (For an introduction to randomly composed automata and their stochastic properties see Malmgren 1984.)

The choice of randomly composed finite automata does not only make for simplicity but also shows the generality of the model in the sense that in it, very few assumptions about specific structure is needed to produce learning. To be sure, the probability of learning in any single randomly composed automaton is often small but "better" data result if large systems of randomly composed automata with cooperative output are used. One way of modelling such cooperative output is to let the most frequent response of the system's subautomata determine the output of the whole system (Malmgren 1984).

Simple adaptation and low-order comparison in randomly composed automata

As a preamble to the discussion of natural resonance I want to point out some trivial but nonetheless important facts about the relation between change of input and change of state in finite deterministic systems.

In an earlier article (Malmgren 1984) I described a simple possible application to biological systems of the notion of randomly composed automata. Such an automaton tends towards a stable state under a constant sequence of inputs. On the other hand, if a new input is given the system usually changes its state. Supposing that an organism contains a large number of such automata, the changes of state of which are directly reflected in the amount of change of the organism's output, the phenomena of sensory adaptation and behavioural habituation immediately receive a plausible explanation. See Malmgren (1984) for details, including calculations and simulation experiments.

If one wants to speculate about the psychology of automata (or, to be more serious: if one wants ultimately to explain psychological notions in terms of mechanistic systems), one might say that a randomly composed automaton "*compares*" the successive inputs and reacts to a "*mismatch*" (over time) by *changing its state*. The typical reaction to a match is instead *no change of state* (or as one might also say, "no reaction at all").

Higher-order input comparisons in automata

Suppose now that a randomly composed N state automaton with input is exposed to a *recurrent input sequence* $I(t_1), I(t_2), \dots, I(t_{k+1})=I(t_1), \dots$ which - for the simplicity of the argument - contains k different inputs which may be renumbered I_1, \dots, I_k . At some time $T \leq kN+1$ it will occur that the automaton is exposed to the same input while being in the same state as it was at some previous time when that particular input occurred. The trajectory of the automaton's state under the repetitive sequence then forms a complex limit cycle C_1 with length $\leq T$. Suppose now that the input sequence is changed so that I_{k+2} is substituted for $I_j, j \leq k+1$. Since the automaton is randomly composed it will most probably switch to another limit cycle C_2 (which may of course have elements in common with C_1). So, again the novelty of the input is reflected as a change in the behaviour of the system.

The switch to another input sequence may imply a switch to a more constant input. This is the case if the sequence is changed from 121212... to 111111... In this case one may say that a global (higher-order) novelty occurs entailing that the input from now on contains less local (lower-order) novelty. It is clear that both these factors will tend to be reflected in the behaviour of a randomly composed automaton in the sense that there will probably be a large-scale change from a closed trajectory including a fair amount of change to a closed trajectory

involving less change. This is because longer trajectories are more probable with the alternating sequence than with the constant one. For example, it is more probable that the automaton goes to a steady state under 111111... than under 121212..., and cycles with length $>N$ are simply not possible under constant input.

The above argument can be generalized to cover changes between other high-level input constraints than the repetitive sequences described above, but for the moment we will confine our attention to this special case. We will now re-describe the facts by saying that the randomly composed automaton simultaneously performs low-level and high-level comparisons and tends to react not only to mismatches between inputs but also to mismatches between successive sequences.

Now, if input changes over time tend to change the state or the trajectory of automata, why not dispense with specifically designed "reset-on-mismatch" mechanisms (as those hypothesized in ART) and instead use novelty as the sole kind of error signal? In the next section I will show how low-level input mismatches over time can be used as externally produced error signals which tend to move the system into a state having a desired output.

The problem which will then remain is of course the following one: how could one possibly dispense with the externally produced error signal if the system only detects mismatches *over time*? After all, the system is expected to learn not only that its present input matches its past input but also that, in some sense, its present response is the correct one - i.e., that its present output matches its present input. And how could it learn that, if no process is designed to perform such a matching?

In the section after the following one, I will explain the basic idea of natural resonance: how, in systems with a certain kind of feed-back, high-level input mismatches over time internally produce an adequate error message which signals a mismatch between present input and output, i.e. a signal which tends to move the system into a closed trajectory in which the system performs correctly.

Novelty as an external error signal

I want to substantiate the present conception of an error signal by referring to an earlier model of operant conditioning, i.e. learning with an external but unspecific teacher mechanism (Malmgren 1985). The hidden ("memory") unit of the system investigated there is a randomly composed finite inert automaton, i.e. a finite automaton with a transition matrix which is randomly constructed except for a rather strong bias towards stability. Usually the memory unit is shielded from environmental information, which is equivalent to the assumption that it receives a constant input (here called the "background" input). This means that it tends strongly toward stable states (longer limit cycles do occur but not as often as in completely random automata). The system's output is a certain combination (by means of a randomly chosen function) of the memory state and the present stimulus, conforming to the model of a modifiable reflex (see Figure 1). The outputs are conceptualized as corresponding to step-wise movements in two dimensions.

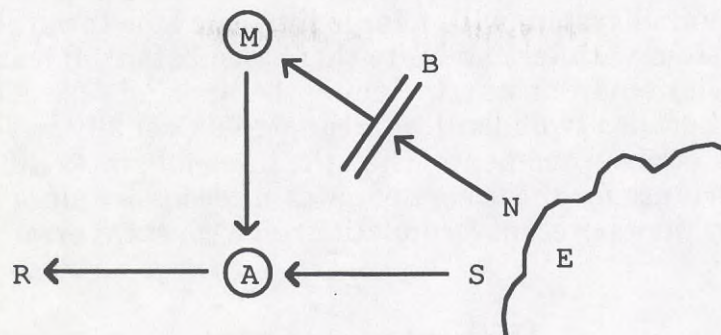


Figure 1. Selection of stable states through temporary intrusion of environmental noise. M = Memory unit; B = Information barrier; N = Noise; E = Environment; A = Reflex apparatus; S = Stimulus; R = Response.

The system moves in a noisy environment (i.e., with many possible stimuli) which contains an operant contingency: if a certain stimulus occurs and the system does not react by moving in a pre-determined direction (one of four), the shield which protects the memory unit from environmental information is temporarily removed - one should think of this as corresponding to the *alerting* of an organism. The signal that is let through (hereafter referred to as the "error" signal) differs from the background input usually received by the memory unit, and so tends to move the unit from one stable state to another. If the system is "lucky" this new state determines a correct response to the next occurrence of the "dangerous" stimulus. Otherwise a new error signal of the same kind occurs.

Over time, the system will perform a kind of random walk among those states which constitute possible attractor states (and limit cycles) under background stimulation, until a "successful" state is reached or the random walk has itself become trapped in an unsuccessful limit cycle. Already in very small inert automata with only 16 possible states, the final count of correct responses amounts to 75-80% (see Malmgren 1985) although the error signal is totally undirected. (For a possible extension of the model including targeted error signals, cf. Malmgren & Östensson 1989.)

For the present discussion it is important to point out that the essence of the error signal in the earlier model is the *novelty* of the signal. At different unsuccessful trials, the signal which is received by the memory unit may be any one of those constituting the noise. Its effect, if any, is due to the simple fact that it is another signal than the usual one. Note that the novelty which plays a role is not only novelty vis-a-vis the background input but also novelty vis-a-vis the previous unsuccessful trial (if there was one). If the error signal is intrinsically the same at two consecutive trials, the chances are lower that a failure will be followed by a success than they are in the case where the two error signals are different. This is of course because in our sense of "compare", the system compares not only successive stimuli but also successive trials.

So, the indeterminacy of the error signal is not only allowed by the model but conducive to its good performance. This holds especially for the case of a small system with a large intrinsic bias towards stability. In a large system with less stability there is substantial learning even with a less noisy environment, including the limiting case of a constant error signal, because two identical error signals usually have different effects if they occur at moments when the internal states differ. This is of some importance for the theory of natural resonance since it depends rather heavily (in its present formulation) on a constant error signal.

Natural resonance

The system unit here consists of two elements, the (inner) (or *memory*) element and the outer (or *resonant*) element. The memory element has an output which we will call the *internal representation*, for reasons that will be clear later. The resonant element combines the external input from the environment (the *stimulus*) with the internal representation and computes a certain kind of function, the resonance function, from them. The result is sent to the memory element. Hence the memory element does not take the environmental stimulus as such as input but something which is a function of this stimulus and the previous output of the memory element itself.

If one accepts the above description in terms of internal representations one can say that the memory unit only perceives the environment through its own interpretation of it. Therefore we will sometimes refer to the result of the computation in the resonant element as the *percept*. In several applications we will however instead think of the output of the resonant element as the *motor* response of the whole unit.




Figure 2. M = Memory element; R = Resonant element; INREP = Internal representation; MP = Perceptual (or motor) response.

One simple kind of resonance function has the following properties. For all external inputs except one, the percept is identical to the external input. For the remaining external input, the *zero stimulus*, the percept is identical to the internal representation. This corresponds to a matrix of the kind described in Table 1.

Table 1
The Computational Matrix of a Simple Resonant Function

Representation	Stimulus					
	S ₀	S ₁	S ₂	S ₃	..	S _M
R ₁	P ₁	P ₁	P ₂	P ₃	..	P _M
R ₂	P ₂	P ₁	P ₂	P ₃	..	P _M
R ₃	P ₃	P ₁	P ₂	P ₃	..	P _M
·	·	·	·	·	..	·
R _M	P _M	P ₁	P ₂	P ₃	..	P _M

To help understanding the point of this construction, one should think of S₀ as a state of "no information about the environment", in which case the internal representation completely determines the response of the resonant element. Similarly, S₁ to S_M correspond to a state of "full information about the environment", in which the external stimulus completely determines the resonant response. It may also be helpful to think of S₀ as being transmitted via some physically independent channel (a gating mechanism), but other interpretations are of course possible.

Note that if the first column is not considered, the matrix has a diagonal such that $P(i,i) = R_i = P(0,i)$. Also, for every representational element R_i, S₀ and S_i are the only stimuli which give the percept P_i. In other words, the i:th representation plus the zero stimulus (= no environmental information) is functionally equivalent to the i:th stimulus. Therefore one can say that the i:th representation *represents* the i:th stimulus - or to be more precise, represents the i:th stimulus under the zero stimulus condition.

The representation relation introduced here is defined in terms of functional relations and does not entail any qualitative similarity between representation and reality. A fortiori, the concepts of *match* and *mismatch* between representation and reality (stimulus) cannot be defined in terms of similarity. To be sure, in computer simulations the easiest way out is to use the numeral 2 as a name for R₂, S₂ and P₂ alike, because then the results become more easily interpretable (cf below). But one should not confuse computational heuristics with the properties of the model itself.

Indeed, the present concept of representation excludes all kinds of simultaneous comparisons (on the part of the system unit) between representation and stimulus, since *when any stimulus except zero is present the representation does not matter at all to the future behaviour of the system*. This is why this conception of a representation is intrinsically connected with comparisons over time.

Simple resonance functions also include all cases where one representation is functionally equivalent under the zero stimulus to a disjunction of other stimuli. Indeed, a great lot of all logically possible $n \times m$ matrixes contain some resonant part, in this sense, and therefore

some rudimentary representation. But we will confine our attention to the special case just described.

A still more general class of resonance functions is created if one takes account of other degrees of environmental determination than one or zero. In continuous resonant elements, a non-linear interaction between stimulus and the output from the memory unit can form the basis of a natural-resonance theory of representation. Non-linear interaction means that for certain values of the external stimulus (say, close to zero) the influence of the internal output (representation) is higher than it is for stimulus values in other regions. In other words, under certain conditions the internal output tends to determine the resonant response, while under other conditions the external stimuli tends to dominate the response. This paves the way for the possibility that a certain internally produced signal may be functionally similar (although perhaps not identical), under the first condition, to a certain class of external stimuli fulfilling the second condition. Then there will be representation in a wide sense. However, in the discrete automaton model considered here we need not take this complication into account.

Adaptation to stimulus sequences

I will now show that the above-mentioned construction will in itself entail a tendency towards learning of high-level stimulus constraints, if these are presented repeatedly and in alternation with zero stimuli.

Suppose a repetitive sequence $S_1S_2\dots S_kS_1S_2\dots$ is presented to the system. After a while it has entered a complex limit cycle C_1 of memory states. During this limit cycle it also produces a repetitive sequence of internal outputs (representations), which however have no effect on the system since it is completely environment-driven.

Now switch the input condition to a sequence of zeroes. In other words, the automaton is given time to *internally rehearse* what happened. It may happen, although it is in general not very probable, that it manages to stay in exactly the same limit cycle. Under zero input the system is completely representation-driven. So if it does not stay in the same cycle but enters another one, C_2 , this must be due to the fact that the sequence of representations which it produces does not "match" the sequence of external stimuli previously given (in the sense of "matching" given above). This in turn entails that the representations produced under the previous constraint did not match the stimuli.

Supposing that the system did not stay in the original limit cycle when "rehearsing" the first time, repeat the stimulation. When the original repetitive condition is again presented, the system may or may not switch to a completely new limit cycle C_3 . If it does, it might happen that the sequence of representations produced "matches" the stimuli. If that is the case (and the match includes phase relations), the behaviour of the system during the next "rehearsal" will be unchanged - it has learnt. If there is a mismatch, there is a possibility that the system will again switch to another limit cycle, and so on.

There is however always a "risk" that the system enters a complex limit cycle of unsuccessful trials involving mismatch, since it may well

happen at any time that the memory unit does not recognize the difference between the stimulus sequence and a representation sequence which does not match it. So learning is not in any sense guaranteed. The argument only shows that under the conditions outlined above, the learned state with a complete match between representations and stimuli does tend to occur at the end of the experiment more often than at the beginning.

In the finite automaton realization which I have simulated, the memory unit is usually a randomly composed automaton with a state-dependent internal output which can take 4 different values (i.e., there are 4 different representations). The system has an external input which can take 5 values including zero and it contains a resonant element of the kind depicted in Fig. 2 and Table 1.

I have exposed 10.000 such 20-state automata (starting in a randomly chosen state) to a regular sequence containing 4 occurrences of sub-sequences consisting of one single stimulus or of three different stimuli, i.e. 1 1 1 1 or 123 123 123 123, alternating with rehearsal periods consisting of 4 and 12 zeroes (to keep phase) respectively. At each rehearsal it is noted whether the automaton's internal representations have reached a stable limit cycle which coincides with the input sequence. The number of correct sub-sequences within each trial is also recorded. After the first 7 trials, the input constraint is changed to 3 3 3 3 and 341 341 341 341 respectively. Finally the automaton is run towards equilibrium for 8 trials (starting in the same original state as before) with the last-mentioned stimulus condition; in all rehearsals the representations proper to the *first* constraint are counted. This last part of the experiment is performed in order to exclude the possibility that the stable solutions are due to an unspecific effect of repetitiveness.

In Table 2, the figures for "correct" refers to the total number of correct responses during a trial of 4 sequences which means (since there are 4 possible representations) that for condition "1", the number expected from random performance is 10.000. The figures for stable solutions are absolute numbers (at most 1 success per trial).

Table 2
Adaptation to Simple and Complex Repetitive Input in Naturally Resonant Automata. For Further Explanation, See Text Condition "1"

Correct responses on 16 rehearsals		Stable solutions on 16 rehearsals	
1	10009	1	377
2	11752	2	1023
3	12291	3	1261
4	12375	4	1327
5	12411	5	1344
6	12431	6	1347
7	12414	7	1350
8	9952	8	491
9	9350	9	269
10	9192	10	291
11	9095	11	281
12	9177	12	304
13	9122	13	283
14	9076	14	285
15	9158	15	295
16	9105	16	289

Condition "123"

Correct responses on 16 rehearsals		Stable solutions on 16 rehearsals	
1	1989	1	64
2	2522	2	233
3	2671	3	294
4	2669	4	305
5	2716	5	309
6	2670	6	309
7	2686	7	309
8	1940	8	77
9	1894	9	41
10	1926	10	48
11	1915	11	51
12	1934	12	50
13	1914	13	52
14	1913	14	49
15	1922	15	51
16	1929	16	51

These figures may not seem very impressive. However there is an approximately threefold and fivefold increase, respectively, of stable correct solutions and there is no doubt that the effects are specific to the training.

With collective output, not only the stable solutions will count. If the most frequent resonant output is allowed to determine the global output of a reasonably large system consisting of randomly composed sub-

automata, such a system will reliably learn to imitate the behaviour induced in it by repetitive sequences.

Classical conditioning

Let us now consider associative learning in the classical sense of Pavlov. I will now show by an analytical argument and by simulations that the resonant automaton will tend to learn a stimulus contingency, and that it shows true conditioning to sequences of stimuli.

In the case of classical conditioning, we may regard the resonant output as being either a percept or a motor response. I here use automata which are random except for the fact that they are rather inert to one of the stimuli, called number 3 and used as the "background" input. This is simply because the effects are otherwise overshadowed by the interstimulus period as a source of variation. Suppose that the system is fed a continuing sequence of 3's which is however sometimes (at randomly chosen moments) interrupted by the short sequence 42 followed by 1. (This is the classical conditioning paradigm with CS=42 and UCS=1.) After a number of such training trials, the system is given the opportunity to rehearse by giving a 0 instead of a 1 for some trials, then the next set of training trials follows, and so on.

Consider one training trial. Since the reaction on the background stimulus 3 is inert, we may simplify the discussion by concentrating on the case where the system enters some stable state during background stimulation. When the first CS-UCS combination has occurred, the system most probably has transited to another stable state. Sooner or later the movement in state space enters a complex limit cycle which may or may not have the property that the representation produced when the UCS occurs is always 1. Now let the system rehearse. If the match between UCS and representation was perfect (as just defined), then the system will not move from its limit cycle, i.e. it will continue to present a "matching" representation during rehearsal. If the match was not perfect, there is a definite probability (depending, among other things, on the proportion of matching representations during the limit cycle) that the system will move to another limit cycle, and so on.

The similarity to the previous case is considerable, but there is an important difference in that phase relations are not important in the case of associative conditioning (since any number of 3's can keep the message).

The amount of learning which is achieved is illustrated by the following simulation. Here 5.000 100-state randomly composed automata with a 50% extra inertia to input 3 were exposed to the conditions just described for 5 periods, each including at least 4 training trials and 4 rehearsals. In addition, in a 6th period the sequence 22 was substituted for 42 as the CS in order to test the hypothesis that the learning actually pertains to the whole sequence and not only to the input 2. In each of the 6 periods it was noted whether or not the internal representation corresponding to 1 occurred at the moment after the CS and whether this was the case for all of the first 4 cycles of the period (most probably indicating a stable solution). Also, the representation

during the last background stimulation (IS in Table 3) was noted in order to detect any generalization to other stimuli than the CS.

As in Table 2, the figures for "correct" in Table 3 refer to the total number of "hits" in 4 samples while stability was of course sampled only once in each period.

Table 3
Classical Conditioning in Naturally Resonant Automata. For Further Explanation, See Text

Period:	1	2	3	4	5	6
CS (correct):	5151	6090	6398	6407	6448	5036
CS ("stable"):	143	603	665	699	712	176
IS (correct):	4768	4897	4960	4887	4998	4897
IS (stable):	80	221	240	234	254	108

It can be seen that the probability of a "correct response" rises gradually to approximately 130% of the value expected from complete randomness. Also, there is a sharp drop to near the base level when the alternative CS (22) is introduced, indicating that the system has learnt that it is the sequence 42 - not the input 2 in any context - which is followed by 1. It seems to me that this model could in principle serve to explain certain naturally occurring conditioning phenomena, supposing cooperative output.

References

- Doya, K. and Yoshizawa, S. (1989). Adaptive neural oscillator using continuous-time back-propagation learning. *Neural Networks*, 2, 375-85.
- Grossberg, S. (1987). *The Adaptive Brain. Vol. I.* Amsterdam: North-Holland.
- MacKay, D.M. (1956). The epistemological problem for automata. In C.E. Shannon and J. McCarthy (Eds.), *Automata Studies*. Princeton, N.J.: Princeton University Press.
- Malmgren, H. (1984). Habituation and associative learning in random mixtures of deterministic automata. *Göteborg Psychological Reports* 14:2.
- Malmgren, H. (1985). On the nature of reinforcement. *Göteborg Psychological Reports* 15:3.
- Malmgren, H. and Östensson, O (1989). Selective attention is selective learning. *Göteborg Psychological Reports* 19:3.



