

**An ART Model
of
Discrimination Learning**

Arjan Berkeljon

UNIVERSITEIT VAN AMSTERDAM
FACULTY OF SOCIAL AND BEHAVIOURAL SCIENCES
DEPARTMENT OF PSYCHOLOGY

An ART Model of Discrimination Learning

by

Arjan Berkeljon
Studentnumber 9906509

A research report (*werkstukverslag*) submitted in
partial fulfillment of the requirements of the degree

Master of Science in Psychology
(*Doctorandus in de Psychologie*)

Supervised by
dr. M.E.J. Raijmakers, dr. I. Visser,
& prof. dr. J.M.J. Murre

UNIVERSITEIT VAN AMSTERDAM
FACULTY OF SOCIAL AND BEHAVIOURAL SCIENCES
DEPARTMENT OF PSYCHOLOGY

Zaandam, April 2006

An ART Model of Discrimination Learning

Abstract

This paper addresses the ability of an ART based neural network to model human quantitative performance and qualitative aspects of the learning process on a discrimination learning task. These qualitative aspects of learning were recently studied by Schmittmann et al. [2005]. An analysis of a data set of human subjects aged 4 to 20 revealed two distinct learning modes in the learning process: (1) a sudden rational learning process by means of hypothesis testing; and (2) a slower but also sudden learning process and not an incremental learning process as has been previously suggested [see Kendler, 1979]. The present study uses an adapted version of an ART network as employed by Levine and Prueitt [1989] in their study of a task related to discrimination learning, the Wisconsin Card Sorting task. Learning in the network is guided by an attentional mechanism which selectively affects current sensory processing based on previous reinforcement. The network is able to model developmental differences through separate adjustment of the valuation of positive and negative reinforcement. Results show that although this network is able to model quantitative human performance through a rather sophisticated learning process, two distinct learning modes cannot be distinguished in its learning behavior. Two structural limitations related to this finding are discussed and possible modifications are suggested to further expand and utilize this network's promising structure.

Contents

Contents	i
1 Introduction	1
1.1 Discrimination learning	1
1.2 An ART model of discrimination learning	4
2 Method	7
2.1 WCST modeling: Levine & Prueitt [1989]	7
2.2 Modeling a discrimination learning task	13
3 Results	18
4 Discussion	23
A Appendix	26
A.1 Network specifications Levine & Prueitt [1989]	26
A.2 Network specifications discrimination learning network	28
References	31

1

Introduction

Studying discrimination learning is an apt method to gain insight into the development of rudimentary category and concept formation. A typical discrimination learning task requires test subjects to choose the correct stimulus from a pair of stimuli. The stimuli presented differ on two or more dimensions (e.g. color and shape). Each of the dimensions has two possible values (e.g. white and black; circle and triangle). Since the task is so simple in nature, subjects of around age 4 and up can be tested using this type of task. This makes it an ideal and often used task in studies concerning the developmental processes in the type of category learning it tests.

Studying the discrimination learning task across different age groups reveals marked differences in performance between them. Such studies commonly use the discrimination-shift task. In this task, an initial stimulus-response relation is learned between a value on a dimension and a response, as described above. After a certain number of correct responses, the reinforcement contingencies are altered (shifted).

The performance differences between adults and children are differences in overall learning speed but also differences in the actual learning process. In this paper these differences will be studied using a neural network model of discrimination learning. The network used includes a modifiable mechanism that is thought to model developmental changes relevant to categorization task of which discrimination learning is a prime case.

1.1 Discrimination learning

This section provides an introduction discrimination learning. Empirical studies as well as neural network studies of discrimination learning are discussed.

Empirical findings: two modes of learning

Any study of discrimination learning must take into account the empirical findings that have been accumulated over the years. Reviews of studies on discrimination learning are given by Wolff [1967], Esposito [1975] and Kendler [1979, 1995]. A robust finding is that adults and children older than 10 years of age find criterion shifts where reinforcement changes such that the opposing value on the same dimension as before is now rewarded the easiest, i.e. subjects

learn them quickly. For example reinforcement might change such that while white was rewarded before, black is now rewarded. These types of shifts are known as *reversal* shifts.

More difficult, and hence slowly learned are *non-reversal* shifts. Here reinforcement changes such that a value on another dimension than before is now rewarded. For example reinforcement might change from black to triangle.

For younger children the results from empirical studies are quite equivocal. Kendler [1979], for example, reports a gradual increase in performance of young children approaching that of adults around age 10. Esposito [1975], however, points out that much variability exists in these studies. Some studies show an advantage for reversal shifts while others show an advantage for non-reversal shifts. This issue centers around the interpretation of younger children's performance on discrimination learning tasks. Ultimately, the question is whether or not children can conceptualize a particular input stimulus as specific instance of a general dimension. It is clear that adults are guided by such mediating conceptualizations, for children this is not the case. The non-systematicity observed in children's performance in at least some studies, certainly suggests that children below a certain age are not guided by mediating conceptualizations but instead employ associative stimulus-response learning.

It is important to note that after a reversal shift responses to *all* of the stimulus pairs need to be changed while after a non-reversal shift responses to only *half* of the stimulus pairs need to be changed. With this in mind a prediction about the different learning processes involved in shift-learning, concept-mediated and associative stimulus-response learning, can be made. With stimulus-response based learning a non-reversal shift should be learned faster than a reversal shift since only half of the stimulus-response relations have changed. With concept-mediated learning however, a reversal shift should be learned faster than a non-reversal shift since only the link between the mediating concept and the response needs to be changed [e.g. Kendler and Kendler, 1962; Raijmakers et al., 1996].

An interesting result regarding the different results obtained for adults and young children is presented Raijmakers et al. [2001]. A finite mixture model analysis¹ conducted on simple discrimination learning data gathered from subjects ages 6 to 10, revealed that a two-component mixture model best fitted the data. The mixture model consists of a component for fast learners and a component for slow learners. These components represent the distribution of errors made in learning. It is believed that the rational-learning component can account for learning based on hypothesis testing and mediation by concept-formation, like in adults. The slow-learning component is thought to account for the generally non-systematic learning present in younger children (ages 6 and below).

The results of Raijmakers et al. share characteristics with the *Levels-of-function* theory [Kendler, 1979]. In this theory, two distinct modes of learning are posited, meant to account for children's equivocal performance. The first mode of learning is an hypothesis-testing, rational, and fast learning process. The second mode of learning is a slow and incremental process. Evidence found by Raijmakers et al. appears to corroborate a dichotomy of learning processes.

¹A statistical modeling technique used to investigate whether a particular distribution is bimodal, i.e. if the distribution is best described as consisting of two component distributions.

It appears, however, that the exact nature of these processes, namely rational versus incremental, is not as Kendler describes.

In a recent study by Schmittmann et al. [2005] performance of 250 subjects ages 4 to 20 on a discrimination learning task was analyzed. Learning was investigated by fitting a number of latent Markov models. In concordance with the above results, we can distinguish two model categories: a model for rational, hypothesis based (fast) learning, and a model for slow learning. Contrary to Kendler [1979] however the model for slow learning also contained an implementation of sudden, as opposed to incremental, learning. The results of the analysis in Schmittmann et al. thus corroborate the finding of two distinct modes of learning, a rational mode and a slow mode. The best fit on the data was obtained using a mixture model containing a rational-learning component and a slow-learning component.

Neural network modeling

In addition to the empirical studies discussed so far, there have also been studies using neural networks to model discrimination learning. These models can be seen as extensions of the empirical knowledge gathered so far. They are designed to mimic human performance as closely as possible and shed light on possible functional processes involved in discrimination learning.

Raijmakers et al. [1996] describe a simulation of discrimination learning using a feedforward PDP network with backpropagation error-correction. This network contained three layers: one layer of eight input nodes for the stimulus dimensions (shape and color); a second, hidden layer of two *or* four nodes (this was varied); and a third layer of two output nodes. In addition to varying the number of hidden layer nodes, the connections between the hidden layer nodes and the input nodes was also varied. In so-called unconstrained networks every hidden layer node was connected to all input nodes. In constrained networks, however, every hidden layer node was only connected to the input nodes of *one* dimension.

Results of this network on a simple discrimination learning task as described above were unambiguous for networks with four hidden layer units. These networks learned a non-reversal shift faster than a reversal shift (unconstrained and constrained). Also, unconstrained networks with two hidden layer units learned a non-reversal shift faster than a reversal shift. This was not so for constrained networks with two hidden layer units. For these networks there was no difference in learning on both shifts. A trial-by-trial analysis showed that this network performed supraoptimal on trials where the stimulus pairs remained the same after a non-reversal shift (recall that for a non-reversal shift only half of the responses change).

These results suggest that learning in this network proceeds by establishing stimulus-response relationships without mediating concept-formation. The authors conclude from this that these types of PDP networks are thus not suitable to model human performance on discrimination learning tasks. In humans, it clearly is the assumption that learning is, at least after a certain age, concept-mediated.

The doubts regarding the suitability of PDP networks for modeling discrimination learning expressed by Raijmakers et al. [1996] are not shared by Sirois and Shultz [1998]. They claim the chosen network topology had too much

of an influence on the results reported by Raijmakers et al.. However, since Raijmakers et al. acknowledge that *certain* neural network models may well offer insight into the process of discrimination learning, the points raised by Sirois and Shultz hardly carry any weight. Fortunately Sirois and Shultz also suggest an implementation to back up their claim that a feedforward network can indeed model human discrimination learning adequately.

Their model is a cascade-correlation network, a minimal feedforward network that can change its own topology during learning. When learning stalls, it recruits hidden layer nodes to track the error in the responses it gives. Thereby it increases its computational ability and problem-solving capacity.

The network used four input nodes and two output nodes. Developmental differences were mediated by adjusting the sensitivity of the network. A highly sensitive network is more sensitive to discrepancies of desired target and network output and thus requires more learning to achieve an adequate output based on a desired target. Less sensitive networks on the other hand, allow for larger discrepancies between desired target and network output and thus require less learning to reach adequate output. Higher sensitivity thus leads to a form of overtraining which in turn increases depth of learning. The authors claim that this overtraining can account for the differences in shift-learning behavior. They claim adult performance is best modeled by highly sensitive networks in which overtraining and thus better learning takes place. Children's performance, on the other hand, is best modeled by networks with lower sensitivity which means learning will be less sophisticated.

Sirois and Shultz report that their high sensitive networks indeed learn a reversal shift faster than a non-reversal shift. For the less sensitive networks there is no such difference. These results seem to support earlier empirical findings (see above). Interestingly enough, the network learned the task without recruiting any mediating hidden layer nodes. The authors take this as support for their claim that no conceptual mediation is necessary to perform a discrimination learning task.

One major caveat with their interpretation is that they did not study the sequential learning process of the model (the progression of correct and incorrect responses). In light of the results obtained by Schmittmann et al. [2005], the ability of an incrementally learning network to adequately model a non-incremental learning process such as discrimination learning, can be doubted. Therefore, the conclusion drawn in Raijmakers et al. [1996], that simple feedforward networks cannot capture all there is to discrimination learning, is not defeated by Sirois and Shultz's arguments.

1.2 An ART model of discrimination learning

To extend the modeling results reported above, and the results reported by Schmittmann et al. [2005], this paper describes an attempt to capture the sequential learning behavior that is apparent in discrimination learning in a neural network model. The approach used here employs a model from Levine and Prueitt [1989] adapted to discrimination learning. This model is built upon the ART 1 structure [Carpenter and Grossberg, 1987].

ART 1 is a neural network architecture based on the *Adaptive Resonance Theory* by Grossberg [1976a,b]. ART networks are self-organizing, adaptive,

and pattern classifying networks with adequate biological and psychological plausibility [Carpenter and Grossberg, 1987; Raijmakers and Molenaar, 1997]. ART 1 classifies binary input patterns which makes it applicable to discrimination learning. Furthermore, the ART architecture seems well-suited to model the sequential nature apparent in many learning processes, and thus discrimination learning as well [Raijmakers et al., 1996]. This is so because ART networks are designed to cope with the stability-plasticity dilemma [Carpenter and Grossberg, 1987, 1988]. This dilemma refers to the problem any learning system faces; on the one hand it has to remain plastic and adaptive to new and significant input; on the other hand, it has to remain stable when faced with irrelevant input. ART networks solve this by using a hypothesis-testing based learning process in which previously learned information gives rise to expectations that focus attention selectively on relevant parts of new input.

In Levine and Prueitt [1989], the ART structure was expanded by adding an attentional mechanism of biases acting on the process of hypothesis-based learning in the ART network. This expansion incorporates a modifiable parameter that controls the strength of this attentional mechanism. Varying this strength allows the network to model the effects of neurobiological damage and developmental processes on cognitive tasks.

The original version of the model, used by Levine and Prueitt [1989], was designed to simulate the effects of lesions in the frontal cortex on performance on the Wisconsin Card Sorting Test (WCST), a learning task related to discrimination learning. In the WCST subjects are required to sort a deck of cards according to one of three rules: color of the objects on the cards, shape of the objects on the cards, or number of objects on the cards (the objects are different geometrical shapes of different colors). Sorting according to the rule (which is unknown to the subject), is brought about by reinforcement learning. After a fixed number of response in accordance with the rule, a new sorting rule is introduced and the sorting resumes.

The WCST is relevant to this study because it of the interesting finding that the mistakes children make in performing this task are similar to the mistakes adults with lesions in frontal cortex make. These mistakes, also known as perseveration errors, consist of the subject being unable to comply with a new sorting rule after one has already been learned. Both frontally lesioned patients and children have trouble making this adjustment. It has been suggested that this error in children is the result of a lack of inhibition of the previously learned rule [Crone et al., 2004]. Taken together, the similarity in errors made by children and frontally lesioned patients, and the capacity of the Levine and Prueitt model to simulate frontal lobe damage by means of a modifiable attentional mechanism, suggest this model is suitable to capture children's performance on the WCST and related tasks.

The current study is an implementation of the Levine and Prueitt [1989] model, adapted to a discrimination learning task. The model will be evaluated by comparing it to the results reported by Schmittmann et al. [2005] using a latent Markov model analysis on the learning sequences produced by the model. The model will be tested on a simple discrimination learning task using color and shape as the dimensions. It is expected that the model will be able to simulate the fast, rational learning process as well as the slow, sudden learning process, both found in the Schmittmann et al. study. By modifying the attentional system of the network, it is expected that the network will be

able to model behavior of subjects at various ages on a discrimination learning task. In other words, it is expected that modifying the attentional system will result in a class of networks where the learning process is best described by a statistical model of rational, fast-learning; and in a class of networks where the learning process is best described by a statistical model of sudden learning.

Each of these issues will be discussed in more detail in the remaining part of this paper. Three main topics will be discussed in the following order:

1. findings of a replication study of Levine and Prueitt [1989] are reported; an implementation and evaluation of the model is presented.
2. an implementation and the workings of the Levine and Prueitt model adapted to discrimination learning is described.
3. performance of this model on a simple discrimination learning task is described and analyzed; results of a latent Markov model analysis are presented and compared to those obtained in Schmittmann et al. using the same method.

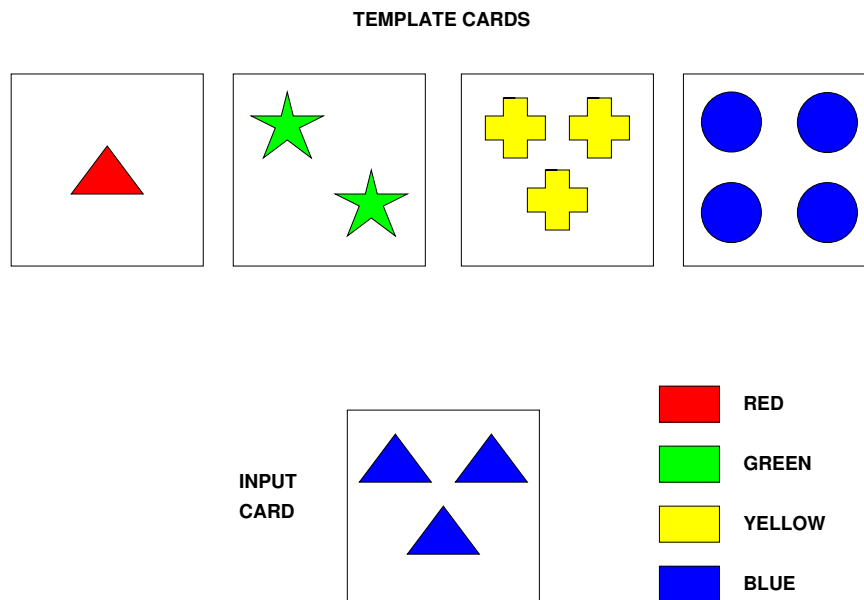
2

Method

2.1 WCST modeling: Levine & Prueitt [1989]

Levine and Prueitt [1989] describe results of a study modeling the effects of frontal lobe lesions on the Wisconsin Card Sorting Task (WCST). This task consist of sorting a deck of cards according to certain sorting rules pertaining to the color, shape and number of objects depicted on the cards. Sorting is done by matching an input card to one of four template cards (see figure 2.1). After a certain number of responses in accordance with the rule, the sorting rule is changed. Sorting should now proceed according to this different criterion.

Figure 2.1: Cards used in the Wisconsin Card Sorting Task [Adapted from Levine & Prueitt, 1989]



Levine and Prueitt modeled the so-called perseveration errors frequently made by frontally lesioned patients (and young children) in this task. These errors consist of a rigid adherence to one particular sorting rule even though it

or answer of the network. If the card chosen and the input card share a feature a so-called match signal is sent to the bias and habit nodes for that feature. This either increases or decreases the activity of the bias node depending on the reinforcement.

Reinforcement is mediated by a parameter α , a simplified implementation of cognitive functions controlled or mediated by the frontal cortex such as guidance of behavior and motivation. This parameter is assumed to be high in normal subjects and low in frontally lesioned patients. The mediating function α has on bias activity thus enables the network to model behavioral effects of frontal lesions in the WCST.

Replicating Levine and Prueitt [1989]

To assess the network's capabilities and structure we set out to replicate Levine and Prueitt [1989]. Exact specifications are supplied in the appendix (A.1). Equations and parameter values are as in Levine and Prueitt unless otherwise indicated.

Initial attempts to replicate Levine and Prueitt's findings were disappointing. The exact implementation is rather underspecified in the article. One notable problem concerned the time the nodes were allowed to update, which is only vaguely mentioned in the article.

This problem was most apparent in the behavior of the bias nodes. Because the bias node activities only reach an equilibrium after a certain amount of time their end-value depends crucially on the amount of time they are allowed to update. With the initial update time chosen (50 time steps for integration according to the equations specified in the appendix), bias performance was erratic. Some nodes responded in an expected manner while others did not.

Figure 2.3(a), for example, shows the responses of two bias nodes; one represents a dimension not in accordance with the sorting rule and the other represents a dimension that is in accordance with the sorting rule.

As figure 2.3(a) shows, both bias nodes start from an initial value and increase and decrease, respectively before finally reaching an equilibrium. This is to be expected, the activation value of the bias node in accordance with the rule increases and the activation value of the bias node not in accordance with the rule decreases.

Unfortunately not all bias nodes exhibited this behavior. A substantial number behaved as the node shown in figure 2.3(b). This node appears to show the correct behavior at first. It represents an incorrect dimension and its activity drops, which is expected. However, shortly thereafter its activity level increases and levels out only slightly below its initial value. This surprising behavior appeared whenever an ambiguous input card was supplied to the network.

Clearly the amount of time the bias nodes are allowed to update can have significant influence on their final value. If the node from figure 2.3(b) is only updated very briefly, its value will be substantially lower than its initial value. This is what is expected of a node representing an incorrect dimension. However arbitrary this may seem, shortening the update time indeed turned out to be the right solution to the problem of erratically performing bias nodes. Instead of 50 time steps, the integration process of the equations describing the

activities of the various nodes in the network was now set very small: only 0.025 time steps.¹

Results

Adjustment of the update time resulted in much better performance. As can be seen in table 2.1 the network used in the current study follows the pattern of results obtained by Levine and Prueitt [1989]. In fact, it even appears to learn the rules after the first one faster than the original network.

With regards to table 2.1 it is important to note that in both studies no fixed input sequence was used. In the WCST, however, it is common to use a fixed input sequence. Levine and Prueitt programmed the input sequence by hand (which was supposed to be fairly random) whereas we used a pseudorandom sequence.

To test for possible effects of the input sequence on the results we also ran 10 simulations with our network using a fixed WCST sequence. The results are given in table 2.2. T-tests revealed no significant differences between any of the reported means in the right column of table 2.1 and those in table 2.2.


Conclusion

Using the same parameter values the network in the current study can qualitatively replicate the results obtained by Levine and Prueitt [1989]. Quantitatively the network used here learned slightly faster than Levine and Prueitt report. One likely explanation for this difference is a difference in the procedure used to update the activation values of the x , y , bias, and habit nodes. Levine and Prueitt used a FORTRAN routine called NONAUT to solve the differential equations used in the network.² The current study used the R extension `odesolve` as differential equation solver. Currently it cannot be assessed what the exact differences between these procedures are. However, since the remaining equations and parameters used in both studies are equal, *some* kind of difference in the differential equation solvers is a likely cause of the observed differences in performance of both networks.

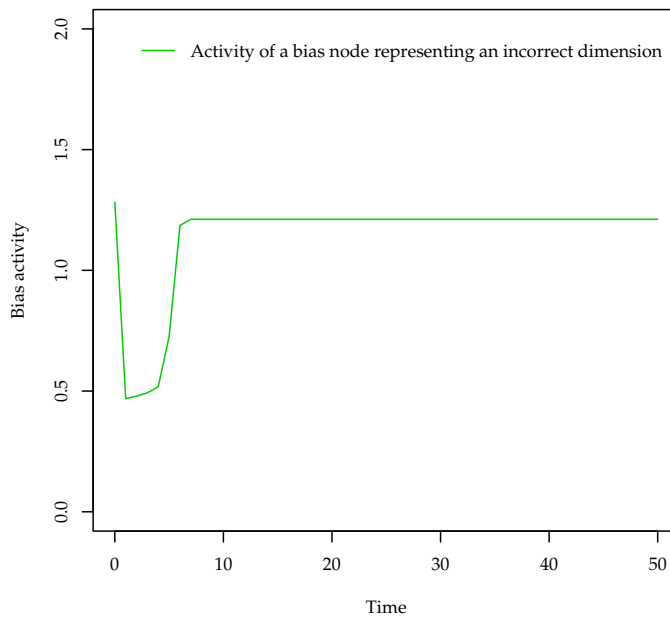
Since the most interesting aspect of the network's performance is its qualitative aspect and not its exact performance, this should pose no problem to further exploration of the network's capacities, however. Even if the network does not exactly replicate the data found in empirical studies of discrimination learning, these results suggest a qualitative comparison should certainly be possible.

¹Clarification in this matter was obtained through one of the original authors, D.S. Levine, in personal communication with Dr. Raijmakers.

²This information was also obtained by Dr. Raijmakers in personal communication with D.S. Levine.



(a) Expected bias node responses



(b) Unexpected bias node response

Figure 2.3: Bias node responses

Table 2.1: Comparison of results WCST simulation

Levine and Prueitt [1989]			Current study		
	Criterion	Trial	Criterion	Trial	SD
$\alpha = 4$ ("Normal")	color	13 ¹	color	13.7 ²	1.42
	shape	40	shape	31.8	5.51
	number	82	number	48.4	10.49
	color (again)	96	color (again)	60.3	10.38
	shape (again)	115	shape (again)	76.8	10.42
				number (again)	97.2
$\alpha = 1.5$ ("Frontally damaged")	color	13	color	13.8	2.86
	Thereafter color		Thereafter color		

¹ Trial refers to the first trial on which the network achieved ten correct matches in a row on the given criterion

² Trial refers to the same as above only now averaged over 10 simulations.

Table 2.2: Results with a fixed input sequence

	Criterion	Trial	SD
$\alpha = 4$ ("Normal")	color	13.9	3.14
	shape	35.9	10.79
	number	53.0	12.80
	color (again)	64.6	13.18
	shape (again)	82.3	12.66
	number (again)	104.1	13.44
$\alpha = 1.5$ ("Frontally damaged")	color	15.2	4.96
	Thereafter color		

Results are the amount of trials to criterion averaged over 10 simulations

2.2 Modeling a discrimination learning task

Adapting Levine and Prueitt [1989]

The network used in Levine and Prueitt [1989] was extended in order to be usable with a discrimination learning task. A number of changes were made (see figure 2.2 and appendix A.2 for exact network specifications).

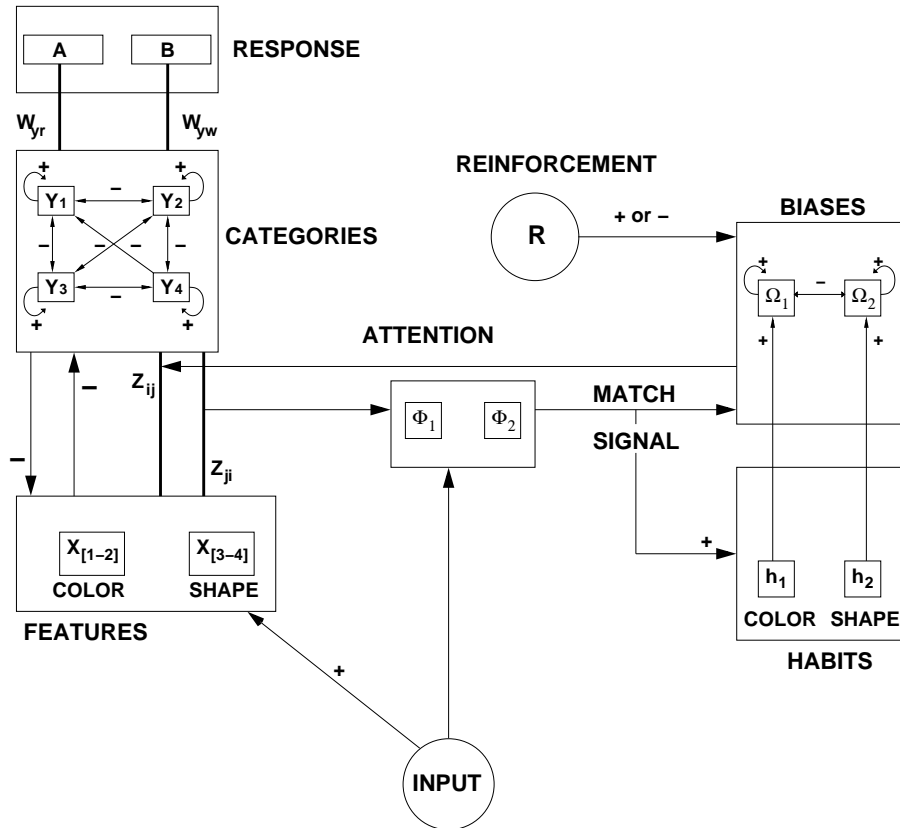


Figure 2.4: Network used to model the discrimination learning task [Adapted from Levine and Prueitt, 1989]

Input to the network is represented one stimulus at a time (an object either white or black and circle or triangle). The stimulus characteristics are coded by four input nodes, one for each possible characteristic, white, black, circle, and triangle. Input propagates to four category nodes, again one for each characteristic. Signals to the category nodes are mediated by the bias nodes, with one node per dimension (color and shape). Activation from the category nodes propagates to the response nodes (A and B), of which one will end up with the highest activation value. This is then taken to represent the choice the network makes. The network is correct if response node A has the highest activation value and the input agrees with the rule. It is also correct if response node B has the highest activation value and the input does not agree with the rule. The weights between the response and category nodes represent an

initial preference for one characteristic over the other on a particular dimension. Dimensional preference is determined by bias nodes and is discussed below.

The category node activation represents the value the network attaches to each of the stimulus characteristics. The category node with the highest activation thus determines the most relevant stimulus characteristic and consequently the most relevant stimulus dimension. This information is used to compute a match signal between category response and input which is sent to the bias nodes. Bias node activation is thus selectively enhanced for that dimension on which the network is currently classifying. This match signal also selectively enhances habit node activity, which represents how often a classification on a particular dimension has been made.

Based on the response, the input and the rule, feedback is either given or withheld. This feedback either increases or decreases bias node activity mediated by two parameters, α^+ and α^- . These function analogously to the α parameter in Levine and Prueitt [1989]. The parameter was split into two components to control for the effects of ambiguous input following a rule change. In the WCST certain cards fit a template card on more than one dimension. Should such a card be presented right after a rule change, learning in a network with only one α parameter will be hindered on this trial. Splitting the α parameter into two different components separates the effects of positive and negative feedback on bias node activity much better. This is especially beneficial in learning a discrimination task because here the input can always be classified on both the color and shape dimension, i.e. input is always ambiguous.

With the α component split in two, the average value of the two components determines overall speed of learning, a lower average resulting in slower learning. The ratio of α^+ to α^- controls the difference between ‘easy learning’ and ‘hard learning’. Easy learning being a situation where the network has a preference in accordance with the current rule (i.e. category to response weights and initial bias activation are both set accordingly) and hard learning being a situation where the network has a preference opposite to the current rule. A high α component ratio, will mainly hinder *hard* learning. *Easy* learning, on the contrary, should remain relatively unaffected.

Simulation methodology

A total of 256 simulations were run divided into two conditions:

1. 128 networks were set to model child-like performance ($\alpha^+ = 1.60$ and $\alpha^- = 0.1$)
2. 128 networks were set to model adult performance ($\alpha^+ = 4$ and $\alpha^- = 22$)

Values of α^+ and α^- were obtained using the original α values (1.5 and 4) from Levine and Prueitt [1989] as a benchmark. Values of α^+ and α^- were calibrated during test simulations to give adequate performance. The present values were found to provide a stable and useful response pattern for the purposes of this study.

The weights between the category and response nodes were set randomly per pair at 0.4 and 0.6 at the beginning of each simulation. This expresses a slight preference for whichever response node has the largest initial activation.

Initial bias node activations were chosen uniformly random from the range 0.98 – 1.02 with the constraint that their average activation is equal to 1. This range was based on benchmark simulations with the bias nodes set to 1. The bias node activations after a criterion was learned were taken to represent a good initial bias value. They were within the 0.98 – 1.02 range.

The simulations were run on an adapted version of a discrimination learning task. Stimuli were presented one at a time and had to be classified by the network in one of two categories represented by the response node with the largest activation value following an input stimulus. Stimuli differed on two dimensions, color and shape. These dimensions each had two possible values: white and black, and triangle and circle. Recall that any input signal is a 4-place vector with value 5 for dimensional characteristics present in the input and 0 for a dimensional feature absent from the input (e.g. the input “a white triangle” is given by the vector (5,0,5,0)). The input sequence was chosen randomly at the beginning of each simulation.

There are four possible rules the network can learn in this version of the task; white is correct; black is correct; triangle is correct; circle is correct. In the current study we did not study shift-behavior so each network had only one rule to learn. In each condition of 128 networks, 32 networks learned each rule. A rule was considered learned if the network gave 10 correct responses in a row. This is a criterion commonly used in discrimination learning studies with human subjects.

Statistics

The overall effect of α on learning (e.g. slow learning with a low average α and faster learning with a high average α) will be assessed by comparing the overall means of the number of trials to criterion between the two conditions, low α and high α , using a T-test. The data within each condition will be analyzed using an analysis of variance on the mean number of trials to criterion in each condition in order to test for a difference between easy and hard learning. As explained above, it is expected that in the low α condition, where the α component ratio is highest, the difference between easy learning and hard learning is most prominent. This should be evidenced by an interaction between the rule to be learned and initial preference (bias). In the high α condition, where the α component ratio is smaller, there should be no or only a small interaction between criterion and initial preference.

Learning models

To test for the existence of distinct modes of learning a Markov models analysis was used [see Visser et al., 2002, for the method]. Two learning models were fitted to the data, a concept-identification (CI) model and an all-or-none (AN) model. Additionally a mixture (CI-AN) model with a concept-identification component and an all-or-none component was fitted to the data. Figure 2.5 shows a graphical representation of the CI and AN model. These models were chosen following Schmittmann et al. [2005].

The CI model is meant to model an hypothesis-based, rational learning process. Initially one hypothesis is chosen randomly (e.g. “White is correct”).

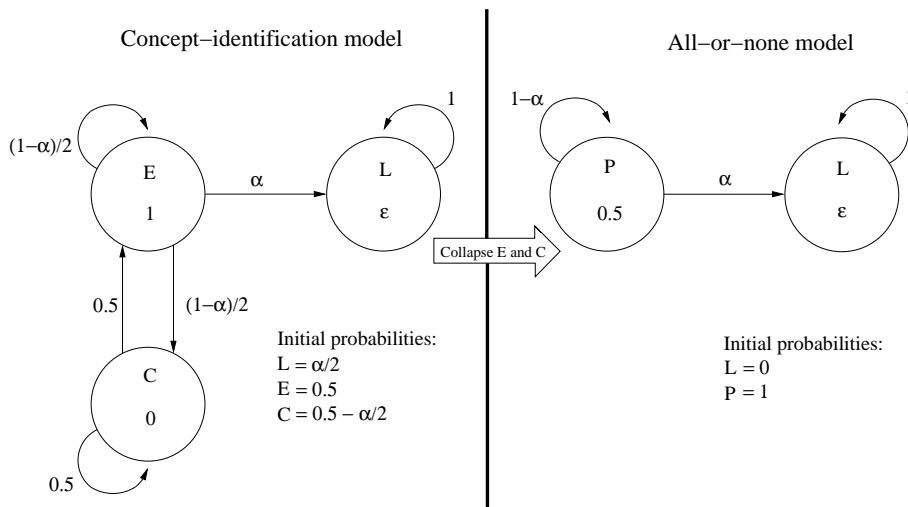


Figure 2.5: Learning models. Circles represent states (L: Learned, E: Error, C: Correct, P: Pre-solution). The number (or ε) in each circle is the probability of an error while in that particular state. Line arrows represent state transitions; probabilities are listed accordingly. The block arrow represent the relation between the two models. [Reprinted from: Schmittmann et al., 2005]

After an error subjects choose from the remaining hypotheses until no more errors are made and the learned state is reached.

The CI model consists of three states: the learned state (L), where the correct hypothesis (rule) is used; and two pre-solution states, E and C, where an incorrect hypothesis (rule) is used and either the incorrect response is given (E), or the correct response is given (C). In this model learning can only occur following an error. Thus the probability of moving from state C to state L is zero. Moving from state E to state L is governed by the learning parameter α .³

Since learning only occurs after an error, and an error implies that now only half of the possible hypotheses remain to be chosen from, the initial probability of the learned state (L) is $\alpha/2$ given that the learning rate is α . Since there are two possible answer choices the initial probability of being in the error state (E) is 0.5. The initial probability of being in the correct response state thus is $0.5 - \alpha/2$.

In state C both moving from C to state E as well as staying in state C have probability 0.5; one can either stick with the current hypothesis or try the alternative. Moving from state E to state C or remaining in state E both have probability $(1 - \alpha)/2$. Moving from state E to state L has probability α , the learning parameter. Once in the learned state subjects will remain in this state. The learning parameter is expected to be 0.5 since after an error two of the four possible hypotheses remain. The probability of an error in the learned state, ε is expected to be 0 since the learning criterion was 10 correct responses in a row. Unlike human subjects, who might make an error even after

³Not to be confused with the α parameter of the neural network model used in this study.

having learned the correct rule, the network used here will apply the same rule without error until the criterion changes.

The AN model is much simpler. It is meant to model an hypothesis-sampling strategy. All hypotheses are equally probable during the entire pre-solution state. In other words there is no (or diminished) feedback processing between trials.

Subjects are assumed to always start in a pre-solution state (P) hence the initial probability is 1. Moving from P to the learned state L has probability α , the learning parameter. Remaining in P has probability $1 - \alpha$. As in the CI model, once in the learning state subjects will remain there. Since the AN is meant to model slow learning, the learning parameter is expected to be much smaller in the AN model than in the CI model. As in the CI model, once the network has learned the correct rule it will apply it consistently. Therefore the probability of an error in the learned state is expected to be 0.

Although the CI and the AN model are qualitatively quite different, they are technically very similar. As can be seen from figure 2.5 collapsing the E and C states of the CI model into one pre-solution state gives the AN model. In this pre-solution state (P) learning can occur after an error and a correct response with the same probability for each.

3

Results

Simulation results

Results of the simulations are shown in table 3.1. Reported are the mean number of trials it took each network to learn each criterion. Rows divide the results by which criterion the networks had to learn. The columns divide the results by initial bias that was given to the network.

A T-test on the overall means of both conditions, low α ($M = 18.12$) and high α ($M = 14.19$), reveals they differ significantly, $t(140.49) = 8.60$, $p < 0.001$. As expected the high α networks learn faster than the low α networks.

In addition two analyses of variance were done on the data for the low α networks and for that data of the high α networks. As expected there is a significant interaction of criterion and bias in the low α condition, $F(1, 124) = 52.78$, $p < 0.001$. The off-diagonal means all differ significantly (Tukey's HSD: $p < 0.001$). As expected there is also a significant interaction of criterion and bias in the high α condition, $F(1, 124) = 29.67$, $p < 0.001$. The off-diagonal means differ significantly (vertical off-diagonals, Tukey's HSD: $p < 0.001$; horizontal off-diagonals, Tukey's HSD: $p < 0.01$).

These data suggest a bimodal distribution of number of trials to criterion for the low α networks and a unimodal distribution of number of trials to criterion for the high α networks.

Markov model analysis

The results of the Markov model analysis are given in table 3.2.¹

Contrary to what is suggested by the results from the analyses of variance, where performance appeared bimodal for the low α group, the best fit was obtained by a single component concept-identification model for both high and low α groups. The best fit is printed in bold in table 3.2. AIC and BIC are criteria for comparing models and a relatively lower AIC or BIC indicates a better fit of the model on the data [see Visser et al., 2002].

Thus, contrary to expectation, the data of the low α group are not best fitted by a CI-AN mixture model with a component for both slow, all-or-none,

¹I would like to thank Verena Schmittmann and Ingmar Visser for help with the Markov model analyses.

Table 3.1: Results of networks with a low α (the ‘children networks’) and a high α (the ‘adult networks’). Means (M) of nr. of trials to criterion, standard deviation (SD), and number of networks per category (N) are reported.

Condition: low α						
<i>(M = 18.12, SD = 5.04, N = 128)</i>						
	Color bias			Shape bias		
Criterion	M	SD	N	M	SD	N
Color	15.00	4.78	28	21.42	3.91	36
Shape	20.26	3.63	27	15.00	4.48	37

Condition: high α						
<i>(M = 14.19, SD = 1.16, N = 128)</i>						
	Color bias			Shape bias		
Criterion	M	SD	N	M	SD	N
Color	13.69	0.78	32	14.75	1.32	32
Shape	14.69	1.27	29	13.71	0.79	35

and fast, concept-identification, learners. The data in the high α group are, as was expected, best fitted by a single component concept-identification model.

Parameter estimates

For the best fitting models, i.e. the unconstrained concept-identification model for both high and low α groups, the learning parameter was estimated. For the low α group the learning parameter was estimated at 0.325. The probability of an error while in the learned state was estimated at 0.042. For the high α group the learning parameter was estimated at 0.396. The probability of an error in the learned state was 0. Recall that the learning parameter in the high α group is was expected to be 0.5 and the learning parameter in the low α group was expected to be smaller than the learning parameter in the high α group. Thus the learning parameter in the high α group is lower than expected while the learning parameter in the low α group is higher than expected, resulting in a minimal difference between them.

Backward learning curves

Backward learning curves can be used to gain insight in a sequential learning process. By aligning the data on the trial of learning and plotting the proportion of errors as a function of time (indicated by trials before or after the trial of learning), the learning process can be graphically represented. The length of the curve indicates the speed of learning, a long curve is indicative of a low learning rate while a shorter curve is indicative of a higher learning rate. The trials before learning is achieved, the pre-solution trials, are especially interesting since these indicate the progression from an unlearned state into a learned state. If this part of the curve is more or less flat, this suggests a stable learning process.

Figure 3.1 shows backward learning curves for both low and high α groups. These curves show the error proportion at each trial of the simulations averaged over simulations. Error proportion ranges from 0 for no error to 1 for all error responses. The dashed line indicates where learning was achieved. Because the curves show the error proportion on each trial of learning the standard error is especially large in the earlier trials. The peak just before learning occurring in both curves is an artifact. Because the criterion for learning was 10 correct responses in a row, it is necessarily so that response before a series of 10 correct response is an incorrect response. Since the data was aligned on the trial of learning (i.e. the beginning of the series of 10 correct response) such a peak is unavoidable using the present method. Averaging error proportion over trial blocks would alleviate this problem but the plotted curves are more informative if error proportion is plotted per trial, hence this approach was chosen.

The curve of the low α groups lies around 0.5. A regression analysis from the fourth trial after learning onset (-16) to the last trial before learning (-1) shows that error proportion starts around 0.5 and then slightly drops ($b = -0.02$, $t(14) = -2.36$, $p < 0.01$). This is expected for the mixture of learning processes this condition was meant to simulate. There are slow 'learners' since the learning process is not as efficient as in the high α group. However, there are also faster 'learners' because when initial preferences are in accordance with the criterion to be learned, learning can be quick in spite of an inefficient learning process. These results are in contrast with the results from the Markov model analysis. As mentioned, a single-component concept-identification model best fitted the data of this group and not, as was expected, a mixture model.

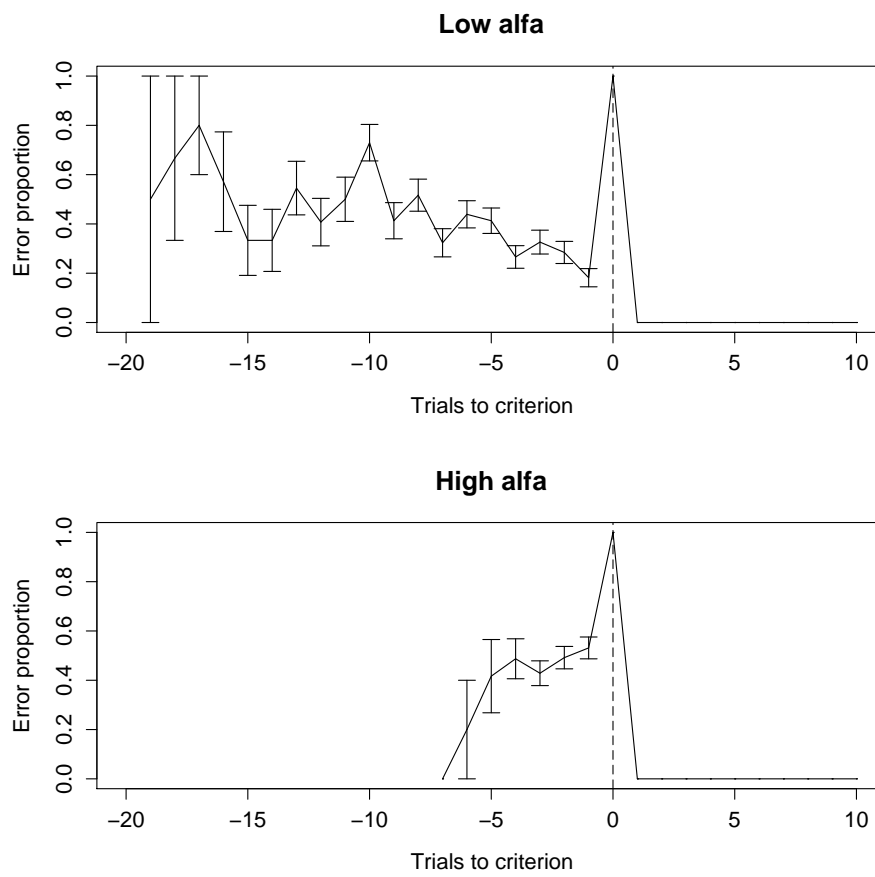
The curve for the high α group differs from the curve of the low α group in that it is shorter, i.e. the high α group learns faster than the low α group. Since the error proportion is very low at learning onset, error proportion rises until the learned state is reached. As with the low α curve, the standard error at the onset of learning is quite large. The slope of the curve is thus inflated by a few networks that generate a correct response early on in the learning process, an effect that is even more pronounced because the curve is short.

Table 3.2: Results of Markov model analyses. Log-likelihood ($\log(L)$), number of free parameters (np), and information criteria (AIC and BIC) are reported.

Low α				
Model	$\log(L)$	np	AIC	BIC
CI	-926.056	2	1856.112	1867.61
CI _{ne}	-1010.275	1	2022.549	2028.298
CI _{free}	-926.018	3	1858.036	1875.283
CI_{nc}	-898.612	5	1807.224	1835.968
AN	-909.798	2	1823.595	1835.093
AN _{ne}	-1033.331	1	2068.662	2074.411
CI-AN	-909.798	5	1829.595	1858.34
CI-AN _{ne}	-926.056	4	1860.112	1883.108
CI _{ne} -AN	-909.798	4	1827.595	1850.591
CI _{ne} -AN _{ne}	-1010,275	3	2026.549	2043.796
CI _{free} -AN	-909.798	6	1831.595	1866.089
CI _{free} -AN _{ne}	-925.691	5	1861.382	1890.126
CI _{nefree} -AN	-909.798	5	1829.595	1858.34
CI _{nefree} -AN _{ne}	-1010.261	4	2028.521	2051.517
High α				
Model	$\log(L)$	np	AIC	BIC
CI	-616.505	2	1237.009	1248.018
CI _{ne}	-616.505	1	1235.009	1240.514
CI _{free}	-614.23	3	1234.459	1250.973
CI_{nc}	-588.148	5	1186.297	1213.819
AN	-603.094	2	1210.188	1221.197
AN _{ne}	-603.094	1	1208.188	1213.693
CI-AN	-603.094	5	1216.188	1243.71
CI-AN _{ne}	-603.094	4	1214.188	1236.206
CI _{ne} -AN	-603.094	4	1214.188	1236.206
CI _{ne} -AN _{ne}	-603.094	3	1212.188	1228.701
CI _{free} -AN	-603.094	6	1218.188	1251.215
CI _{free} -AN _{ne}	-603.094	5	1216.188	1243.71
CI _{nefree} -AN	-603.094	5	1216.188	1243.71
CI _{nefree} -AN _{ne}	-603.094	4	1214.188	1236.206

Note: CI refers to the concept-identification model; AN refers to the all-or-none model; CI-AN refers to the CI-AN mixture model; *ne* refers to models where the probability of an error in state L was set to 0 (no error); and *free* refers to models where the probability of starting in the error state was freely estimated instead of being set to 0.5; *nc* refers to models where all constraints were freely estimated (no constraints).

Figure 3.1: Backward learning curves for low and high α groups. Graphs represent the proportion of errors per trial. Vertical bars represent the standard error. Data was aligned on trial of learning, indicated by a dashed line.



4

Discussion

The two major components of this study were: (1) a replication study of Levine and Prueitt [1989], modeling the effects of frontal lobe lesions on the Wisconsin Card Sorting task using an adapted ART network; and (2) a design and test an implementation of this ART model adapted to discrimination learning using the results obtained in Schmittmann et al. [2005] as a benchmark.

Replicating Levine and Prueitt [1989]

The results obtained by Levine and Prueitt were replicated using an implementation of their network structure rewritten for the present study (see appendices A.1 and A.2). The network used learned the WCST slightly faster than Levine and Prueitt report, however the qualitative pattern of the results was similar. It was suggested that a different method of updating node activation (most notably the bias nodes) is the most likely explanation for the difference in learning speed that was found.

A more fundamental problem, concerning the node activation update routine in Levine and Prueitt and consequently the current study, is the short update period that was chosen for the nodes in network (the input, category, and bias and habit nodes) to update their activation following an input. Contrary to what one would expect from an ART network, the nodes in the network were not allowed to update until an equilibrium state was reached. Instead updating activation was cut off after a certain point.

The choice for this cutoff point remains unmotivated in Levine and Prueitt, and thus similarly so in the present study, save for pragmatic considerations of network performance. A sensible extension of this paper would be to address this issue. It is believed this can be done while preserving the network topology, and thus its defining characteristic of an attentional subsystem with explicitly mediated reinforcement. This would entail adapting/rewriting the differential equations used to describe node behavior to incorporate the requirement of updating to an equilibrium state. An ART implementation where this requirement is met can be found in Raijmakers and Molenaar [1997].

Modeling discrimination-shift learning

A comparison of means of the high α and low α groups shows a significant difference between the two groups in number of trials it took for the network

to learn the criterion. The high α networks learned significantly faster than the low α networks. A comparison of means within the two groups showed that in the low α group, and to a much lesser extent, in the high α group, there was significant interaction between the criterion to be learned and initial preference of the network for one dimension (color or shape) over the other. So far these results are uncontroversial. The low α networks, meant to model children's performance, learn slower than the high α networks, meant to model adult's performance. This is to be expected.

Although the reported results suggest the presence of two distinct groups of 'learners' in the low α group, a Markov model analysis did not yield a best fit for a mixture model. In both the high α and the low α group a single-component concept-identification model (unconstrained) best fit the data.

In the high α group this is certainly an expected result supporting the hypothesis that a high α models a fast, hypothesis-based learning process. For the low α group the result of Markov model analysis is surprising. The hypothesis was that a low α would model a slow learning process but that because of initial preferences slow *and* fast 'learners' would exist (e.g. an slow 'learner' that has an initial preference in accordance with a criterion and is still able to learn quickly). Unfortunately this hypothesis is not supported by the data.

The most natural conclusion is that the learning process in the low α condition does not differ qualitatively from the learning process in the high α condition, it simply differs quantitatively. A large difference between the split α parameters (see section 2.2) simply results in networks that learn slower, not networks that learn differently in a qualitative manner.

More than likely the networks used here all learn according to the same learning process. The difference the α parameter (or split α parameters) makes, is a difference in learning speed. This conclusion is supported by all the results reported in section 3: (1) a Markov model analysis shows that the same single-component CI model fits the data from both high and low α groups; (2) the differences between the groups are predominantly differences in learning speed as evidenced by a comparison of mean trial on which learning was achieved and estimated learning parameter. It took the low α networks longer to learn a criterion than the high α networks and the learning parameter was lower in the low α condition than in the high α condition.

Similarly the difference between the two components of the α parameter mainly appears to mediate the learning speed and not a qualitative difference in the learning process. Recall that the difference between the two components in the high α condition was much smaller than the difference between the two components in the low α condition. Only in the low α condition was there a significant difference in learning speed related to a difference between these two α components. As with the difference between a high and a low α between the two network groups, a difference between α components within a single network group appears to result in a difference in learning speed only and not a difference in learning mode.

As for why the best fitting model was a single-component concept-identification model and not another model it must be noted that in this particular model all parameters were freely estimated. Such a model will fit the data better than the other models simply because there are less constraints (parameters) to be satisfied. Also, one of the main assumptions in the CI

model is that learning can only occur after an error. This was also implicit in one part of the network architecture, namely the weights between the category nodes and the response nodes (see figure 2.2). The weights between the nodes coding the color and shape categories and the responses were only allowed to change after an error (recall that bias and habit nodes did learn continuously and not only after an error). This simplifies network structure, and therefore implementation, but it would be interesting to see the effect of a continuous learning mechanism on these weights.

Prospects

Although the results of this study fail to meet the expectations set out at the beginning, the network structure used here shows promise. It is apparent that the explicitly implemented learning parameter α can be effectively used to model different learning rates. Thus, while the network was able to model fast as well as slow learning, and thus model empirical findings adequately, the differences between these types of learning was a quantitative difference and not a qualitative one.

Even in the absence of a difference between the learning process of the two kinds of networks, i.e. low and high α , the finding that the learning process for both kinds of networks was best captured by a concept-identification model is interesting. Notwithstanding the caveats mentioned above, the fitted model is a proper concept-identification model. This means that the learning process for this neural network model is quite sophisticated and not simple associative learning one would expect from a run-of-the-mill neural network model. This is an important result and points to the potential of the type of network used here.

It is suggested that a possible future implementation of this model should address the problem with the update method of node activations such that the nodes are allowed to update until an equilibrium state is reached. Quite possibly this adjustment will allow for a better separation of the α parameter into two components so that a mixture of learning processes can be modeled. It is also suggested that the learning mechanism for the weights between the category nodes and response nodes is implemented as continuous learning rather than ‘winner-takes-all’ learning as was done in the present study. Hopefully this increased separation of the learning parameter and increased variability in learning will result in a better dissociation of a fast and slow learning processes and quite possibly reveal a qualitative difference between them.

A

Appendix

A.1 Network specifications Levine & Prueitt [1989]

This network (see figure 2.2) was used in Levine and Prueitt [1989] to model the effects of frontal lobe damage on performance on the WCST.

Input to the network is represented by a one-dimensional array $I_{i=1:12}$; $I_{1:4}$ represent the numbers in that order; $I_{5:8}$ represent the colors, red, green, yellow and blue respectively; $I_{9:12}$ represent the shapes, triangle, star, cross and circle respectively. The input, I_i was set to 5 whenever a feature i was present in the stimulus. Otherwise it was set at 0.

Activation of the *feature* nodes $x_{i=1,2,3,4}$ is determined by the input activation, I_i and the *category* node activation, $y_{j=1,2,3,4}$ weighted by z_{ji} . The change in x is given by the following equation:¹

$$\begin{aligned} dx_i/dt = & -Ax_i + (B - Cx_i) \left(I_i + \sum_{j=1}^4 f(y_j)z_{ji} \right) \\ & -Dx_i \sum_{j=1}^4 f(y_j), i = 1, 2, \dots, 12, \end{aligned} \quad (\text{A.1})$$

where the constants are defined as $A = 10$, $B = 5$, $C = 1$, $D = 1$ and f is given by:

$$f(x) = \arctan(x - 1) + \pi/2. \quad (\text{A.2})$$

The input propagates through the *feature* nodes to the *category* nodes and determines the activation of $y_{j=1,2,3,4}$ (each node represents one of four template cards used in the WCST) weighted by z_{ij} and the activation of the *bias* nodes $\Omega_{k=1,2,3}$ (where Ω_1 represents a number bias, Ω_2 a color bias and Ω_3 a shape bias). Thus for each feature node x_i , the corresponding bias number is $k = [(i + 3)/4]$ where $[w]$ is the greatest integer not exceeding w . The change

¹Integration on the equations is done using the *R* extension `odesolve` with timestep 0.025.

in y is given by:

$$\begin{aligned} dy_j/dt = & -Ay_j + (B - Cy_j)(f(y_j) + \sum_{i=1}^{12} g(\Omega_{[(i+3)/4]}x_i)z_{i,j}) \\ & -Dy_j \left(\sum_{r \neq j} f(y_r) + I \right), j = 1, 2, 3, 4, \end{aligned} \quad (\text{A.3})$$

I is a constant of value 100 and g is given by:

$$\begin{aligned} g(x) = & 0, x < .5 \\ & x - .5, .5 \leq x \leq 3 \\ & 2.5, x > 3. \end{aligned} \quad (\text{A.4})$$

The weights z_{ij} and z_{ji} between x and y are fixed. They are large when x and y share a feature and small where x and y are dissimilar. Thus $z_{ij} = 5$ if $i = j$ and $z_{ij} = 0$ otherwise; moreover $z_{ji} = z_{ij}/5$ for all i and j .

The *category* node with the largest activation determines the network's response; it has matched that particular template card to the input given.

Activation of the *bias* nodes is determined by the *bias* node itself, the activation of the corresponding *habit* node and the reinforcement value. This gives the following:

$$\begin{aligned} d\Omega_k/dt = & -E\Omega_k + \left\{ (f - \Omega_k)((h_k - \theta_1)^+ + \alpha R^+ + g(\Omega_k)) \right. \\ & \left. - \Omega_k(\alpha R^- + G \sum_{r \neq k} g(\Omega_r)) \right\} f(\Phi_k), \end{aligned} \quad (\text{A.5})$$

where the constants used are $E = .01$, $F = 3$, $G = 10$ and $\theta_1 = 1$, respectively. The reinforcement value (R) is 1 for a correct response and -1 for an incorrect response. The parameter α is 4 for networks simulating adult behavior on the learning task and 1.5 for networks simulating children's behavior. Variable Φ_k is a match signal occurring between the input and a y node. It is given by:

$$\Phi_k = \sum_{i=4k-3}^{4k} z_{Ji}I_i, \quad (\text{A.6})$$

where i is the index of the x node in the input; J is the index of the most active y node; and k is the index of the corresponding *bias* node. I_i is the input signal (5 or 0).

Activation of the *habit* nodes is given by:

$$dh_k/dt = Hh_k\{(J - h_k)(\Phi_k - \theta_2)^+ - (\Phi_k - \theta_2)^-\}, \quad (\text{A.7})$$

where $H = .1$, $J = 3$ and $\theta_2 = .5$.

The network is partly implemented in R and partly in C ; the network structure and simulation code is written in R while a set of differential equations used to describe the activity of the *feature*, *category*, *bias* and *habit* nodes (see figure 1) were written in C . This part of the code is linked to the R code through an R extension named `odesolve`.

A.2 Network specifications discrimination learning network

The network (2.2) used to model a discrimination learning task is modelled after Levine and Prueitt [1989].

Input to the network is represented by a one-dimensional array $I_{i=1,2,3,4}$; $I_{1,2}$ represent the colors, white or black; $I_{3,4}$ represent the shapes, triangle or square. As in Levine and Prueitt [1989] the input, I_i was set to 5 whenever a feature i was present in the stimulus. Otherwise it was set to 0.

Activation of the *feature* nodes $x_{i=1,2,3,4}$ is determined by the input activation, I_i and the *category* node activation, $y_{j=1,2,3,4}$ weighted by z_{ji} . The change in x is given by the following equation:²

$$\begin{aligned} dx_i/dt = & -Ax_i + (B - Cx_i) \left(I_i + \sum_{j=1}^4 f(y_j)z_{ji} \right) \\ & -Dx_i \sum_{j=1}^4 f(y_j), i = 1, 2, 3, 4, \end{aligned} \quad (\text{A.8})$$

where the constants are defined as $A = 10$, $B = 5$, $C = 1$, $D = 1$ and f is given by:

$$f(x) = \arctan(x - 1) + \pi/2. \quad (\text{A.9})$$

The input propagates through the *feature* nodes to the *category* nodes and determines the activation of $y_{j=1,2,3,4}$ (1=white, 2=black, 3=triangle and 4=square) weighted by z_{ij} and the activation of the *bias* nodes $\Omega_{k=1,2}$ (where Ω_1 represents a color bias and Ω_2 a shape bias). Thus for each feature node x_i , the corresponding bias number is $k = [(i + 1)/2]$ where $[w]$ is the greatest integer not exceeding w . The change in y is given by:

$$\begin{aligned} dy_j/dt = & -Ay_j + (B - Cy_j)(f(y_j) + \sum_{i=1}^4 g(\Omega_{[(i+3)/4]}x_i)z_{i,j}) \\ & -Dy_j \left(\sum_{r \neq j} f(y_r) + I \right), j = 1, 2, 3, 4, \end{aligned} \quad (\text{A.10})$$

I is a constant of value 100 and g is given by:

$$\begin{aligned} g(x) = & 0, x < .5 \\ & x - .5, .5 \leq x \leq 3 \\ & 2.5, x > 3. \end{aligned} \quad (\text{A.11})$$

²Integration on the equations is done using the R extension `odesolve` with timestep 0.025.

The weights z_{ij} and z_{ji} between x and y are fixed. They are large when x and y share a feature and small where x and y are dissimilar. Thus $z_{ij} = 5$ if $i = j$ and $z_{ij} = 0$ otherwise; moreover $z_{ji} = z_{ij}/5$ for all i and j .

The *category* node with the greatest activation determines the network's response. This response is weighted by: $w_{yr,j=1,2,3,4}$, the weights between the y nodes and the category A response node, resp_1 ; and $w_{yw,j=1,2,3,4}$, the weights between the y nodes and the category B response node, resp_2 .

Response activation is given by the following equation [Kruschke, 1996, see].³

$$a^{\text{resp}_{1:2}} = \sum_{j=1}^4 \text{ychoice}_j w_{yr:w,j} \quad (\text{A.12})$$

The variable ychoice represents the activity of the y nodes; $\text{ychoice}_j = 1$ when y_j is the most active y node and 0 otherwise (a 'winner-takes-all' principle). The weights w_{yr} and w_{yw} are randomly initialized as either 0.4 or 0.6, i.e. their average was 0.5 while one was high and the other low. This expresses a slight preference for the highest response nodes. They were updated using the following equations:

$$\Delta w_{yr:w,j} = (t_{r:w,j} - a^{\text{resp}_{1:2}}) \text{ychoice}_j, \quad (\text{A.13})$$

where the target variable $t_r : w$ is set to 1 if the network response is correct and 0 if incorrect. A correct response is classifying in category A when the input pattern accords with the current rule *or* when classifying in category B when the input pattern does not accord with the current rule.

Activation of the *bias* nodes is determined by the *bias* node itself, the activation of the corresponding *habit* node and the reinforcement value. This gives the following:

$$d\Omega_k/dt = -E\Omega_k + \left\{ (f - \Omega_k)((h_k - \theta_1)^+ + \alpha^+ R^+ + g(\Omega_k)) \right. \\ \left. - \Omega_k(\alpha^- R^- + G \sum_{r \neq k} g(\Omega_r)) \right\} f(\Phi_k), \quad (\text{A.14})$$

where the constants used are $E = .01$, $F = 3$, $G = 10$ and $\theta_1 = 1$, respectively. The reinforcement value (R) is 1 for a correct response and -1 for an incorrect response. Parameter α^+ and α^- are 4 and 22 respectively for networks simulating adult behavior on the learning task and 1.60 and 0.1 respectively for networks simulating children's behavior.

Variable Φ_k is a match signal occurring between the input and a y node. It is given by:

$$\Phi_k = \sum_{i=4k-3}^{4k} z_{ji} I_i, \quad (\text{A.15})$$

³Integration of these equations is not done in `odesolve` but in an R function in the simulation code.

where i is the index of the x node in the input; J is the index of the most active y node; and k is the index of the corresponding *bias* node. I_i is the input signal (5 or 0).

Activation of the *habit* nodes is given by:

$$dh_k/dt = Hh_k\{(J - h_k)(\Phi_k - \theta_2)^+ - (\Phi_k - \theta_2)^-\}, \quad (\text{A.16})$$

where $H = .1$, $J = 3$ and $\theta_2 = .5$.

The network is partly implemented in R and partly in C ; the network structure and simulation code is written in R while a set of differential equations used to describe the activity of the *feature*, *category*, *bias* and *habit* nodes (see figure 1) were written in C . This part of the code is linked to the R code through an R extension named `odesolve`.

References

- CARPENTER, G.A., & GROSSBERG, S. [1987]. A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer, Vision, Graphics and Image Processing*, **37**:54–115.
- CARPENTER, G.A., & GROSSBERG, S. [1988]. The ART of adaptive pattern recognition by a self-organizing neural network. *Computer*, **21**:77–88.
- CRONE, E.A., RIDDERINKHOF, K.R., WORM, M., SOMSEN, R.J.M., & VAN DER MOLEN, M.W. [2004]. Switching between spatial stimulus-response mappings: a developmental study of cognitive flexibility. *Developmental Science*, **7**(4):443–455.
- ESPOSITO, N.J. [1975]. Review of discrimination shift learning in young children. *Psychological Bulletin*, **82**:432–455.
- GROSSBERG, S. [1976a]. Adaptive pattern classification and universal recoding I: Parallel development and coding of neural feature detectors. *Biological Cybernetics*, **23**:121–134.
- GROSSBERG, S. [1976b]. Adaptive pattern classification and universal recoding II: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, **23**:187–202.
- KENDLER, H.H., & KENDLER, T.S. [1962]. Vertical and horizontal processes in problem solving. *Psychological Review*, **62**:1–16.
- KENDLER, T.S. [1979]. The development of discrimination learning: A levels-of-functioning explanation. *Advances in Child Development and Behavior*, **13**:83–117.
- KENDLER, T.S. [1995]. *Levels of cognitive development*. Erlbaum, Mahwah, NJ.
- KRUSCHKE, J.K. [1996]. Dimensional relevance shifts in category learning. *Connection Science*, **8**:201–223.
- LEVINE, D.S., & PRUEITT, P.S. [1989]. Modeling some effects of frontal lobe damage—Novelty and perseveration. *Neural networks*, **2**:103–116.
- RAIJMAKERS, M.E.J., & MOLENAAR, P.C.M. [1997]. Exact ART: A complete implementation of an ART network. *Neural networks*, **10**(4):649–669.

- RAIJMAKERS, M.E.J., VAN KOTEN, S., & MOLENAAR, P.C.M. [1996]. On the validity of simulating stagewise development by means of PDP networks: Application of catastrophe analysis and an experimental test of rule-like network performance. *Cognitive Science*, **20**:101–136.
- RAIJMAKERS, M.E.J., DOLAN, C.V., & MOLENAAR, P.C.M. [2001]. Finite mixture distribution models of simple discrimination learning. *Memory & Cognition*, **29**(5):659–677.
- SCHMITTMANN, V.D., VISSER, I., & RAIJMAKERS, M.E.J. [2005]. Multiple learning modes in the development of performance on a rule-based category-learning task. Submitted to *Elsevier Science*.
- SIROIS, S., & SHULTZ, T.R. [1998]. Neural network modeling of developmental effects in discrimination shifts. *Journal of Experimental Child Psychology*, **71**: 235–274.
- VISSER, I., RAIJMAKERS, M.E.J., & MOLENAAR, P.C.M. [2002]. Fitting hidden Markov models to psychological data. *Scientific Programming*, **10**: 185–199.
- WOLFF, J.L. [1967]. Concept-shift and discrimination-reversal learning in humans. *Psychological Bulletin*, **68**(6):369–408.