

# Connectionist simulation of attitude learning: Asymmetries in the acquisition of positive and negative evaluations.

J. Richard Eiser *University of Sheffield*

Russell H. Fazio *Ohio State University*

Tom Stafford, Tony J. Prescott *University of Sheffield*

## Abstract

Connectionist computer simulation was employed to explore the notion that, if attitudes guide approach and avoidance behaviors, false negative beliefs are likely to remain uncorrected for longer than false positive beliefs. In Study 1, following standard backpropagation of error training, a three-layer network with a state unit reflecting cumulative discrimination success (energy) perfectly discriminated good and bad inputs distributed across a two-dimensional space. Under contingent feedback, connection weights were only updated when the network produced outputs representing approach behavior. In this condition, bad inputs were still discriminated but some good inputs were not recognized as such. In Study 2, the network architecture distinguished a system for learning evaluations from a mechanism for selecting actions. Biasing action selection toward approach even when energy was high eliminated the asymmetry between learning of good and bad inputs under contingent feedback. We discuss implications for various attitudinal phenomena and biases in social cognition.

*Keywords:* *Keywords attitude, connectionism, learning, simulation*

This is a preprint of:

Eiser, J. R., Fazio, R. H., Stafford, T., Prescott, T. J. (2003). Connectionist simulation of attitude learning: Asymmetries in the acquisition of positive and negative evaluations. *Personality and Social Psychology Bulletin*, 29(10), 1221-1235.

It is based on the final MS submitted to the publisher and so should be identical except for formatting changes and, additionally, in that figure 1 appears as the authors intended. Please feel free to circulate or cite.

**Correspondence to:** J. Richard Eiser, Centre for Research in Social Attitudes, Department of Psychology University of Sheffield Sheffield S10 2TP, UK. **Email:** j.r.eiser@shef.ac.uk

**Acknowledgements:** This research was supported by a grant R000223077 from the Economic and Social Research Council, UK to the first author and by Senior Scientist Award MH01646 and grant MH38832 from the National Institute of Mental Health to the second author. The assistance of Christopher Watkins with preliminary phases to the simulations is gratefully acknowledged.

## 1 Introduction

Although the concept of attitude has remained central to social psychology for as long as the discipline has existed, there has been remarkably little research directly concerned with how attitudes are acquired. Despite Allport's (1935, p.810) famous definition of attitude as a mental and neural state of readiness, organized through experience, the role that experience actually plays in attitude organization has remained comparatively under-researched compared with other topics, such as the influence of attitudes on behavior. Although textbooks refer frequently to attitude formation and change, the studies described under this heading tell us far more about how we can be influenced to change our existing attitudes than about how we come to have such attitudes in the first place.

Despite this emphasis in the literature, several major theoretical approaches contain explicit or implicit assumptions about the kinds of learning processes that may underlie the acquisition of attitudes. Several models of attitude-behavior relations (e.g., Ajzen, 1991) view behavior as guided by acquired expectancies (Edwards, 1954; Tolman, 1959). More recently, the concept of associative memory is central to a number of models looking at the cognitive and behavioral consequences of attitudes (Fazio, 1990, 1995; Petty & Krosnick, 1995). According to Fazio (1995), attitudes are object-evaluation associations, specifically implying that attitude formation depends on processes of associative learning.

This widespread incorporation of learning notions at a theoretical level, however, has not been matched by much recent empirical work on attitude learning. An early impetus had been provided by Hildum and Brown (1956) and Insko (1965) using operant conditioning notions and by Staats and Staats (1958) within a classical conditioning paradigm. This research, however, ran out of steam, partly because of concerns over participants awareness of the reinforcement or associative contingencies (Page, 1974), and possibly because, for many social psychologists at the time, learning theory paradigms appeared redolent of an outdated behaviorism and incompatible with the current cognitive *Zeitgeist*. Recent social cognition research, however, increasingly acknowledges the importance of priming and other automatic processes occurring below the level of conscious awareness (Bargh, 1997; Fazio, 2001; Wegner, 1994). Classical conditioning of attitudes without explicit detection of covariation has also recently been demonstrated by Olson and Fazio (2001). We therefore believe that the time is ripe for a renewed analysis of the learning processes underlying the acquisition of attitudes.

A priority for such an analysis is to specify the assumptions about learning implicit in more general notions. The idea that we acquire attitudes through associating objects with good and bad experiences is intuitively plausible. For this idea to be the basis of a theory, however, we need to be able to say more precisely how such associations may be formed under different conditions. In particular, we need to ask if there is anything that distinguishes attitude learning from associative learning in general. In other words, is there anything special about learning associations between objects and evaluations as distinct from associations between objects and any other kind of event?

Part of the difficulty with the concept of association is that it can imply a process based merely on the co-occurrence of events. Attitudes can be acquired through passive exposure (see, e.g., Fazio & Zanna, 1981; Olson & Fazio, 2001). However, we decided here to examine the intuition that many of our attitudes may be developed through active exploration of our environment. In such cases, we learn whether we like or dislike different objects, activities or individuals through interacting directly with them. Through such learning, we will choose to engage in activities we find enjoyable and avoid unpleasant activities as far as possible. Our attitudes are shaped by experience, but our attitudes can guide our exploration and hence shape our experience. Attitude learning thus involves a dynamic interaction with the environment, in which our attitudes both guide approach and avoidance behaviors, and are updated by the feedback that such exploration provides. In short, the process underlying the acquisition of such object-evaluation associations may be better defined as a form of reinforcement learning, whereby evaluations are dependent on feedback from the environment, but no feedback is received unless the environment is explored.

Much of this is implied in earlier perspectives on attitude structure, psychological development and group processes by such as Lewin (1936) and Heider (1946). Within Lewin's scheme, goal-oriented behavior is guided by individuals psychological field or life space at a given time, including perceptions of social relationships, while learning (including through social interaction) can lead to restructuring, re-evaluation and differentiation of this life space over time. In a similar vein, Heider's theory of cognitive balance predicts reciprocal interdependence between perceived social relationships and perceived agreement and disagreement, so that, according to the theory, mutual liking increases mutual agreement, and mutual agreement increases mutual liking.

All this implies an asymmetry between how we acquire positive and negative attitudes toward other people and valued objects – at least within a non-coercive environment in which our approach and avoidance behaviors are guided by our expectations of outcome contingencies. In a simple application of the Law of Effect, we will attempt to repeat pleasant experiences and avoid unpleasant ones. Other things being equal, therefore, we should have more experience of positively valued than negatively valued objects (cf. Parducci, 1984). This suggests, however, that most negative attitudes will be more weakly grounded in direct experiential learning than positive ones. Yet at the same time, the evidence from animal learning (Solomon & Wynne, 1954) is that avoidance responses can be very resistant to extinction. For example, a rat that has learnt to move to a different end of its cage when a tone sounds in order to avoid an electric shock will continue to do so, even if the schedule has been changed and no more shocks are given. A cognitive interpretation of this effect is that avoidance prevents the rat from ever experiencing the absence of a shock following a tone in that part of the cage. By contrast, if the rat received a shock in a part of the cage previously experienced as safe, its previous tendency to approach this part would relatively quickly be inhibited. Extended to the human context, this implies that we are more vulnerable to error in our negative attitudes than positive ones. In other words, if we hold a positive attitude toward an object, we will be more likely to approach it and hence have our expectation confirmed or corrected

by experience. However, if we hold a mistakenly negative attitude toward an object and consistently avoid it, we will never, except by chance, discover what we are missing.

### 1.1 Aims of computer simulation

The present paper attempts to explore these ideas through the methodology of connectionist computer simulation. It is important to declare at the outset what the use of such methodology can, and cannot, be expected to achieve. First and foremost, simulation is a technique for clarifying theoretical concepts and predictions. It allows us to ask the question: If such-and-such assumptions are correct, what would be predicted under such-and-such conditions? All forms of theory-building and hypothesis-generation in psychology do this, but computer simulation allows us to do so more precisely, especially when dealing with complex situations and processes evolving over time. Indeed, simulation not merely allows us to be more precise, but demands it of us. In other words, we are forced to define our theoretical assumptions in precise and internally coherent terms, or the program will simply fail.

None of this, by itself, establishes that our theoretical assumptions (or the way we have just specified them) are correct. Simulation is not a substitute for empirical evidence in that sense. However, it helps specify what our assumptions imply, including for situations which would be difficult and/or expensive to reproduce in the laboratory. An example of this is the simulation of social influence processes in groups consisting of many members; the predictions of Heider's (1946) balance theory have been simulated by Eiser, Claessen and Loose (1998) and those of Latané's (1981) social impact theory by Nowak, Szamrej and Latané (1990). Such simulations do not prove, or even confirm the respective theories. However, they help to clarify the theories and extend our understanding of what the theories predict.

The second advantage of simulation is almost the mirror image of the first. All computer simulation depends on a number of inbuilt assumptions, procedural decisions and parameter settings. Many of these find their way into programs as arbitrary or ad hoc fixes, but many others are based on principled theoretical positions concerning the kind of processes being simulated. Even to talk about simulating processes is to take a subtly different position from the view that modeling consists of constructing a rule-based or algorithmic system for finding solutions to problems. Connectionism (e.g. Ellis & Humphreys, 1999; Gurney, 1997; McLeod, Plunkett & Rolls, 1998) refers to a branch of cognitive science that seeks to understand (and simulate) cognitive processes on the basis of rather particular, but simple, assumptions about how knowledge is acquired, organized, retained and recalled within complex systems crudely analogous to a natural brain. Approached in this way, computer simulation can offer process explanations beyond those contained in more traditional theories.

### 1.2 Connectionist principles

The basic idea of connectionism is that brains consist of huge numbers of neurons that can receive information (in the form of electrical activation of varying strength) and then pass this information on to

other neurons. Each neuron, however, is a relatively simple device — essentially a conductor or switch — but brains achieve immense complexity at the system level through the essentially infinite number of ways different sets of neurons can become interconnected with each other. The extent to which information or activation can pass from one neuron to another is determined by the strength of the (synaptic) connection between them.

In order to simulate such processes, connectionist modeling employs systems of simple nodes or units, crudely analogous to (sets of) neurons, that are interconnected to form networks. Different network 'architectures' constrain the way the units pass information to each other. For instance, three-layer nets (as employed here), incorporate a layer set of input units, analogous to sensory receptors, whose levels of activation directly encode the stimuli presented. These input activations are then transmitted to hidden units, that form condensed representations of the input patterns. These then pass activation on to output units, representing the response of the system. The activation of each unit is a function of the sum of the activations it received from other units, weighted by the strengths of the connections to it from each of these units. These connection weights can be either of positive or negative sign, i.e. facilitatory or inhibitory. Simulations involve 'training' the net by adjustment of the connection weights, commonly so as to minimize the discrepancy between the outputs and some designated target values.

Underlying these procedures is the assumption that learning, i.e. experience of covariation between events and feedback from the environment about the goodness of fit or mismatch between predictions and outcomes, results in different patterns of interconnectivity. Connectionist systems thus have no need of a distinct 'memory store'. All the information acquired through learning is stored in the connection weights. These same weights control processes such as categorical perception (Harnad, 1987), generalization and recall of information from partial cues (Hopfield, 1982). This can be expressed by saying that there is no distinction between memory and cognitive processing in connectionist systems (McLeod, et al., 1998).

At least in the hands of many of its practitioners, therefore, connectionism offers, not simply a set of modeling techniques, but an unapologetically theoretical perspective on cognition and learning (McLeod et al., 1998; Seidenberg, 1993). This is potentially generalizable to several topic areas, and not least to social psychology (Eiser, 1994; Nowak, Vallacher, Tesser, & Borkowski, 2000; Read & Miller, 1998; Shultz & Lepper, 1996; Smith, 1996; Smith & DeCoster, 2000). Connectionist simulations do not simply require us to be more precise in specifying our assumptions about the processes underlying, say, balance or social impact theory. Instead, we face the question of whether the core assumptions of our social psychological theory can be translated into, and reformulated within, the more general and often more parsimonious conceptual language of connectionist learning. In this respect, a simulation can be viewed as an invitation to consider principles of connectionist learning as a plausible general theoretical account of the phenomena in question.

### 1.3 Supervised learning and reinforcement learning

There is an important distinction in connectionist theory between supervised learning systems that receive full feedback from the environment as to the correctness of the output generated for each and every input, and reinforcement learning systems that receive only partial feedback. Feedback may be partial in several respects, however, most importantly for the current study, reinforcement feedback is generally contingent on the learning system's outputs interpreted as actions in the environment. If a learning system needs to act in the world in order to receive feedback, then it also needs to explore the alternative actions available in different contexts in order to observe what consequences ensue. A crucial issue for reinforcement learning then, that does not arise for supervised learning, is the need to find an appropriate balance between the exploration of alternative actions, in order to make better future choices, and the exploitation of existing knowledge about feedback contingencies in order to achieve good immediate outcomes. This issue is often termed the exploration/exploitation trade-off (Sutton & Barto, 1998).

Current approaches in reinforcement learning have developed through the combination of connectionist learning techniques with optimal control theory. As this work has progressed, identities, or close similarities, have been discovered between these advances in machine learning, and psychological theories of the processes underlying classical conditioning, instrumental conditioning, and learning with delayed rewards. Since attitude formation can be understood as a process of associative learning that is dependent on exploration and contingent feedback, then it should also be amenable to analysis using this powerful theoretical framework.

### 1.4 Aims of the study

The aim of this paper is therefore to examine, within the methodological and theoretical framework of connectionism and reinforcement learning, possible asymmetries in the acquisition of positive and negative beliefs. We assume that asymmetries may arise from the fact that positive attitudes lead to approach behaviors and hence to increased experience of the attitude object, whereas negative attitudes lead to avoidance and hence less direct experience. We hypothesize that, since such experience is necessary for a learner to build up a correct representation of the environment, false-negative errors (i.e. assuming a good object is bad) will tend to be more frequent than false-positives. Since effective learning requires exploration, we further hypothesize that manipulating the willingness of the system to deviate from its currently preferred action (its exploration strategy), will directly affect its experience of negative objects, and consequently impact on the accuracy of acquired attitudes, and on its overall effectiveness as a learner.

In order to address the problem of choosing actions during learning, it is useful to distinguish two elements of a system that learns from contingent feedback: (i) the learning system, which encodes the

currently preferred actions of the system for any context, and (ii) the action selection mechanism, that chooses an action, at any given time, based on both the currently preferred action (for the current context) and the current exploration strategy. In the paper, we compare two classes of models. In Study 1, the action selected is determined completely by the learning system. In Study 2, the action selection mechanism is separate from, and only probabilistically guided by, the learning system. The performance of the network is influenced by its cumulative record of correct and incorrect responses (here termed energy), but differently in the two studies. In Study 1, energy is connected to the learning system, whereas in Study 2 it only affects the action selection mechanism.

The simulations to be reported attempt to replicate a paradigm developed by Fazio and Eiser (2000) with human participants. In this paradigm, participants are introduced to a computer game in which their task is to survive in a virtual world consisting entirely of beans. Some of these beans are good and provide energy. Others are bad, and eating them results in a loss of energy. Participants are told that if they eat too many bad beans rather than good beans (or go too long without eating at all), they will die. To survive, they need to learn to identify and eat enough good beans, while avoiding the bad beans. The game involves participants being presented with different beans one at a time, and having to choose whether or not to eat them. The beans themselves vary along two attribute dimensions in terms of number of speckles and shape (circular to oblong). There are 10 potential levels on each of these dimensions. Hence, the space of all possible objects comprises a  $10 \times 10$  matrix. The beans actually presented fill 36 of the 100 possible attribute combinations. These are arranged in six blocks or regions, three containing good beans and three containing bad beans (see Figure 1). The main finding from the human data is that participants are much less accurate at identifying good than bad beans. In other words, false negative errors predominate over false positives (where positive means responding to a bean as good). We therefore examine whether an equivalent asymmetry could be reproduced by a neural network constrained to learn only through exploration.

Reproducing the human data, however, is not the sole purpose of the simulations. In the human work, the game is such that feedback is obtained only through approach. But, unlike a net, human learning could be affected by attentional and rehearsal mechanisms. Recent reviews (Baumeister et al., 2001; Rozin & Royzman, 2001) illustrate various senses in which bad is stronger than good. The contribution of the connectionist simulations is that we can examine learning through exploration in a much purer fashion, i.e, unconfounded by any other tendencies that humans might bring to bear on this situation.

Speckles (Y)										
10										
9										
8										
7										
6										
5										
4										
3										
2										
1										
Shape (X)	1	2	3	4	5	6	7	8	9	10

Figure 1: Matrix of input patterns used during training. Clear squares (regions 1, 3 and 5) represent good beans and dark grey squares (regions 2, 4 and 6) represent bad beans.



## 2 Study 1

### 2.1 Method

#### 2.1.1 Network architecture

The code for the simulations to be reported was written specially using the programming software Matlab<sup>TM</sup>, version 5.3.

#### 2.1.2 The learning system

The learning system employed in Study 1 is a fully-connected, three-layer, feed-forward neural network as shown in Figure 2. The input layer of the network comprises 22 units, of which 11 are used to encode one dimension (e.g. shape) and the remaining 11 the other dimension (speckles). These input units take values between 0 and 1, with each level of an attribute being represented by a pattern of activation ( $> 0$ ) across up to six of the 11 units. For example, one speckle would be encoded by the vector  $[1, 1, 0.5, 0.25, 0, 0, 0, 0, 0, 0, 0]$ , four speckles as  $[0, 0.25, 0.5, 1, 1, 0.5, 0.25, 0, 0, 0, 0]$ , through to ten speckles as  $[0, 0, 0, 0, 0, 0, 0, 0.25, 0.5, 1, 1]$ . The effect of this is that any two adjacent levels of an attribute will share one input unit in common where the activation level is at its maximum (1). Because each attribute level is encoded by more than one input unit, and the individual input units contribute to the encoding of more than one attribute level, the network achieves a distributed (rather than localist) representation of the different stimuli. This enables the network to encode location in the space in such a way as to also take account of proximity. Furthermore, because of the roughly Gaussian distribution of lesser activations to either side of the maxima, stimuli up to five steps away from each other on any attribute would share at least one input unit with activation levels  $> 0$ . This was intended to facilitate the generalization of learning to untrained regions of the space. The second layer of the network includes three hidden units, each of which receives activations from all 22 input units. The number of hidden units (reflecting the computational capacity of the network) was determined on the basis of preliminary modeling to be the minimum sufficient for learning these sets of inputs. Also providing input (in Study 1 only) to the three hidden units is a single state unit, effectively a record of the level of energy (i.e. the effects of eating different beans within the context of the game) at a given point in time. The activation of this energy unit varies between 0 and 1, starting the simulation at 1. There is a steady decay in the activation of the energy unit, at the rate of 0.0001 per bean presentation. Eating a good bean increases the activation of the energy unit by 0.001 and eating a bad bean decreases it by 0.001 (whereas *avoiding* a bad bean produces only the time-related decay of 0.0001).

Weighted activations from the three hidden units and the energy unit are then fed through a logistic (squashing) function, restricting the output activation to a range between 0 and 1, and thence to a single output unit. The input to the logistic function from the energy unit is not modified by learning, but

determined by the following hunger function selected on the basis of preliminary modeling, the effect of which is to provide a maximum input of 1 to the logistic function when energy is 0, declining to 0.33 when energy is 1. The effect of this hunger function is that the network is more inclined to eat when energy is low, even if it has not clearly categorized a presented input as a good bean.

$$Hunger = 1/[\gamma \cdot (Energy^2 + Energy + 1)], \text{ with } \gamma = 1. \quad (1)$$

### 2.1.3 The action selection mechanism

In Study 1, the output of the learning system completely determines the action selected. A threshold parameter on the output unit is set at 0.5, at or above which outputs are treated as equivalent to eating a bean, and below which outputs are treated as equivalent to avoidance.

### 2.1.4 Training procedure

The network was trained using variants of the standard backpropagation of error algorithm (Rumelhart, Hinton & Williams, 1986) to modify the connection weights. (Parameter settings were 0.02 for the learning rate and 0.06 for momentum). This form of training requires the target values of the stimuli, i.e. their actual valences, to be defined in advance by the researcher. This allows an error value (or delta,  $\Delta$ ) to be calculated. For this purpose, the output generated by the net in response to a given input is subtracted from the target value for that input. Where the input corresponded to a good bean, the target value was set at 0.9, where it corresponded to a bad bean, it was set at 0.1. A positive  $\Delta$  thus represents an outcome better than the net's prediction, a negative  $\Delta$ , an outcome worse than predicted. This  $\Delta$  is used to modify the connection weights of the hidden-to-output, energy state-to-hidden and input-to-hidden links. The effect of these modifications is to increase the output activation produced in response to a given input if  $\Delta$  is positive, and reduce it if  $\Delta$  is negative.

### 2.1.5 Supervised learning with full feedback

Three variants of this training procedure were used. In the simulations using full feedback,  $\Delta$  was calculated and connection weights were modified both when the net chose to eat (i.e. produced an output equal or greater than 0.5) and when it did not. In other words, the situation is conceptually equivalent to being asked to predict whether a bean was good or bad, and then being simply told whether this prediction was right or wrong. These simulations are an example of a standard discrimination learning problem of the kind that has been the subject of extensive connectionist modeling using supervised learning (e.g. McClelland & Rumelhart, 1988) and provide a benchmark to show how well the network can learn the inputs with no restrictions




Figure 2: Network used in Study 1. Solid arrows indicate connections modified by learning.

### 2.1.6 Reinforcement learning with contingent feedback

More interesting theoretically are the simulations using contingent feedback. In these, connection weights are only modified if the output activation is equal to, or greater than, 0.5 – i.e. if the net has chosen to eat. If the net produces an output of less than 0.5, this is conceptually equivalent to avoiding a bean, and thus receiving no feedback about whether it would have been good or bad. In other words, no learning (modification of weights) takes place on any trial where the output is less than threshold. When this occurs, the net just proceeds to the next input pattern with no modification to the weights. Since feedback is contingent on the action performed, this constitutes a form of reinforcement learning.

### 2.1.7 Reinforcement learning with contingent feedback and confirmation bias.

A further variant of the learning procedure was based on the observation, in the animal learning literature, that avoidance responses appear more resistant to extinction than would be predicted if the non-occurrence of an expected shock following avoidance was processed simply as a non-event. One interpretation (Solomon & Wynne, 1954) is that avoidance is reinforced by a reduction in fear consequential on performance of the avoidance response.

To simulate this form of confirmation bias, we adapted the contingent feedback procedure as follows: on all trials where the network avoided (i.e. produced an output  $< 0.5$ ), *regardless of the true target value for the input*, we calculated a  $\Delta$  as though the target value was 0.1 (i.e. as though it was a bad bean). This  $\Delta$  was then multiplied by an attenuation parameter arbitrarily set at 0.1, and the connection weights were then updated by the back-propagation algorithm in the normal way. In other words, regardless of whether the input in fact corresponded to a good or bad bean, avoidance was reinforced, the strength of the reinforcement (i.e.  $\Delta$ ) being equivalent to one-tenth of that received for correct avoidance under full feedback training for an output activation at the same level. It was deemed necessary to set this parameter relatively low, since otherwise there would be no distinction between correct and incorrect avoidance in terms of their impact on learning.

In both full and contingent feedback conditions, the network was trained for 5000 epochs with all 36 input patterns (beans) being presented once in each epoch, and all weights being updated together (by batch training) at the end of each epoch. For each simulation, the starting state of the network was defined by setting all connections to random values within the range from 0.3 to +0.3. (There were no restrictions on the values taken by connection weights after training). Ten independent replications of each simulation were performed with different sets of initial random weights, analogous to running an experiment with ten independent participants. At the end of training, the output activations corresponding to each of the 36 training input patterns (beans) were recorded. In addition, following training, the network was presented with novel inputs or beans with attribute combinations not previously shown, and output activations were again recorded. The purpose of these test phase trials was to see how the network's representation

of the problem space (instantiated in the connection weights) would produce generalization of learning, i.e. allow input patterns not previously presented to be categorized as good or bad. In this phase, the network was presented with inputs corresponding to all remaining 64 cells of the  $10 \times 10$  matrix, so as to provide complete mappings of the problem space. No  $\Delta$  was calculated on any of these test trials, so there was no further modification of the connection weights during the test phase.

### 3 Results

We first inspected the output activations at the end of 5000 epochs of training in which all 36 input patterns (beans) were presented. For each of the ten replications within each of the three feedback conditions, we calculated (a) the number of correct choices out of 18 for the good and bad beans separately (i.e. outputs of 0.5 or more to good beans and less than 0.5 for bad beans); and (b) the mean (absolute) error ( $\Delta$ ) over the two sets of beans separately (i.e. the differences, regardless of sign, between the output activations achieved after training and the correct target values, averaged over the beans within the good and bad sets). Table 1 shows the means over the 10 replications within each condition. In the full feedback condition, 18 correct responses were obtained to each set of inputs. This was expected from previous connectionist simulations of two-category learning (e.g. McClelland & Rumelhart, 1988). In the other two conditions, the bad beans were all correctly avoided apart from within a single replication under contingent feedback (13 out of 18 avoided). However, 29% of the good beans were also avoided, i.e. categorized as bad. Wilcoxon signed ranks tests indicated that the difference between number of correct choices for good and bad beans was significant under both contingent feedback ( $z = 2.94, p < .005$ ) and confirmation bias ( $z = 2.85, p < .005$ ).

Table 1: **Mean number of correct choices out of eighteen and mean absolute error ( $\Delta$ ) to good and bad input patterns, and mean evaluation of untrained patterns, as a function of feedback (Study 1).**

	N Correct Choices		Mean Absolute Error		Mean Evaluation
	Good	Bad	Good	Bad	Untrained
Full feedback	18.0	18.0	0.02	0.01	0.53
Contingent feedback	12.6	17.5	0.29	0.25	0.42
Confirmation bias	13.1	18.0	0.26	0.09	0.41

The mean absolute error ( $\Delta$ ) scores (i.e. discrepancies from target values) provide more details of the level of learning achieved. Under (unbiased) contingent feedback, although bad beans are consistently avoided,  $\Delta$  remains quite high, and only marginally ( $z = 1.89, p < .06$ ) below that for good beans. This

indicates that, once the outputs to these patterns fell below threshold, they then showed little further polarization toward the correct target value of 0.1. The reason for this is that no further updating of weights then occurred on these trials, and any further improvement in discrimination could only occur as a result of backpropagation of error on the remaining inputs categorized as good. Under confirmation bias, however, the difference between the mean error scores to good and bad beans is more reliable ( $z = 2.20, p < .05$ ), indicating that the output activations for correctly avoided beans continued to decrease toward their true target values. These data were submitted to a  $3 \times 2$  (Feedback  $\times$  Valence) analysis of variance, with repeated measures on the second factor. This revealed a significant effect for Feedback ( $F(2, 27) = 122.45, p < .001$ ), with both the contingent feedback and confirmation bias conditions differing significantly ( $p < .001$ ) from full feedback. The effect of Valence ( $F(2, 27) = 10.02, p < .001$ ) and the interaction ( $F(2, 27) = 4.03, p < .05$ ) were also significant.

We next tested generalization to input patterns not presented during training. Table 1 also shows the average outputs over all 64 untrained test patterns. The three feedback conditions also differed significantly ( $F(2, 27) = 5.32, p < .05$ ), indicating that the untrained patterns were evaluated more negatively under contingent feedback and confirmation bias. In other words, in those conditions where good beans were less well learnt, the network showed a generalization effect, so that novel beans tended, on average to be categorized as bad.

A more complete picture of how the networks generalized from the presented inputs can be seen in Figure 3. This shows the landscapes of mean output activations produced by the network to all 100 possible input patterns in the different conditions, using a monochrome gradation with lighter shades representing higher activations, i.e. more positive valence. For comparison purposes, the top left panel shows the actual target values. The full feedback condition produced a reasonably accurate landscape with hills and valleys corresponding respectively to the good and bad regions, the main difference from the actual target landscape being the spreading (i.e. generalization) of higher and low outputs into the neutral untrained regions. Under contingent feedback, however, the landscapes display incomplete recognition of good beans (particularly region 5), with confirmation bias leading to stronger rejection of regions 4 and 6.

## 4 Study 2

The simulations in the first study demonstrate important differences between learning under full and contingent feedback. Specifically, for the artificial system described, objects that are in fact good (i.e. approachable) may continue to be categorized as bad, and hence be avoided, under conditions where avoidance prevents corrective learning from taking place. Furthermore, in the confirmation bias condition, designed to simulate the presumed reinforcing effect of avoiding an object believed to be bad, negative evaluations of such objects (wrongly) categorized as bad become even more extreme.




Figure 3: Target values (untrained patterns shown as 0.5) and mean output activations for all attribute combinations as a function of feedback (Study 1).

In any simulation study, the effects observed can depend on specific features of the procedure, including, for example, the network architecture and the various parameter settings. (Of course, experimental findings can be just as dependent on procedural details, but this dependency may be less transparent where the methodology is less familiar). Although varying parameters arbitrarily is uninformative, two (related) features of our original network appear particularly relevant to conceptual issues. The first is the relationship between the learning system and the action selection mechanism. In Study 1, the relationship was deterministic, in that the action selected depended entirely on whether the activation of the output of the learning system was above or below threshold. Hence, there was no possibility of exploration except in the context of a positive evaluation. The inclusion of the energy unit, together with the function that computed hunger from its activation, was intended to encourage exploration when energy was low. However, hunger was still located within the learning system in that it provided inputs to both the hidden units and the threshold function on the output unit.

Arguably, a network architecture that distinguished between evaluation (or attitude) and action would be more appealing in terms of its intuitive resemblance to human decision-making. To address this issue, we modified the architecture for Study 2 to differentiate the action selection mechanism explicitly from the learning system, and made the relationship between the two parts of the network probabilistic by adding an element of random noise to the action selection mechanism. The effect of this is that the network had a non-zero probability of approaching beans provisionally categorized as bad, as well as a non-zero probability of avoiding beans provisionally categorized as good.

The second important modification concerns the role of energy or hunger. The energy unit was introduced to allow for the intuition that individuals may be more likely to approach objects about which they may be uncertain, if their temporary need to do so is greater. In short, one is more ready to eat if one is hungry. But here we have an ambiguity. Does hunger make one *think* of a particular food as more appetizing, or merely more ready to *eat* food that one would not normally regard as enjoyable or even edible? In Study 1, the linkage of the energy unit to the learning system at least partly appears to simulate the first of these interpretations. In Study 2, the energy unit was completely separated from the learning system, being linked only to the action selection mechanism. This also enabled us to vary how hunger was calculated from energy, so as to investigate the impact of hunger on exploratory behavior without it having any effect on evaluative learning (i.e. attitude).

## 4.1 Method

Our second study therefore employs an architecture in which the (continuous) evaluation output from the learning system produces a probability distribution from which approach or avoidance behavior is chosen by the action selection mechanism. The energy unit was now treated as part of the action selection mechanism rather than the learning system (See Figure 4). In addition, we varied the specific function whereby the hunger input to the action selection unit was computed from the energy level at any given



time. The basic function was as follows:

$$Hunger = ((Baseline - 1).Energy + 1)^\gamma, \text{ with } \gamma = 5. \quad (2)$$

with the added constraint that the minimum level of Hunger was 0.

Three functions were used as shown in Figure 5. In all three, an energy level of 0 produces a hunger value of 1. In the *neutral* condition ( $Baseline = 0$ ), hunger approaches an asymptote of 0 at  $energy = 1$ . In the *cautious* condition ( $Baseline = -0.6$ ), hunger falls more quickly to 0, with the effect that hunger can only override a negative expectancy when energy drops very low. Finally, in the *risky* condition ( $Baseline = 0.6$ ), the network is prepared to take more risks (i.e. eat beans expected to be somewhat bad) when its energy level is high. The biases produced by these hunger functions were added to the output from the learning system. The resulting judgment (i.e. the evaluation plus hunger) was transformed using a logistic function to produce a probability of eating between 0 and 1. This means that as the judgment of a bean moves further away from neutral (0.5), the probability of approaching the bean rapidly comes to be close to 0 or to be close to 1, depending on the direction in which the judgment diverges from 0.5. Beans with judgments of 0.5, such as might result early in learning with no contribution from the hunger function, have a probability of 50% of being eaten.

Finally, a stochastic probability function was introduced by comparing this result to a randomly generated number between 0 and 1 (labeled noise in Figure 4). If the output exceeded this random number, the action selected would be eat, otherwise avoid. As a bean's evaluation falls it will, because of the stochastic nature of the action selection process, still be eaten occasionally, although with less and less frequency as the evaluation moves closer to 0. Conversely, positively evaluated beans may still be occasionally avoided, though less frequently as their evaluation (plus hunger) approaches 1. Hence, this additional random element weakens the deterministic link between evaluation and action present in Study 1 by allowing the network occasionally to explore some beans toward which the provisional attitude is negative.

In all other respects, the simulations were the same as in Study 1. That is, the same set of input patterns was used, the same form of input coding, and the same algorithm to update the weights in the learning system, under the same three feedback conditions (full, contingent, confirmation bias).

## 4.2 Results

Table 2 shows the mean scores for the number of correct choices,  $\Delta$ s and evaluations of untrained patterns. Output plots (omitting full feedback conditions) are shown in Figure 6. The results of Study 1 were broadly replicated under the neutral and cautious hunger conditions, with the bad beans being even better learnt than before (note the zero values for mean  $\Delta$ s). In other words, the full feedback condition resulted in perfect learning of both good and bad beans, whereas good beans (particularly, as in Study 1, those in region 5) were imperfectly learnt under contingent feedback and confirmation bias. Furthermore,




Figure 4: Network used in Study 2. Solid arrows indicate connections modified by learning.

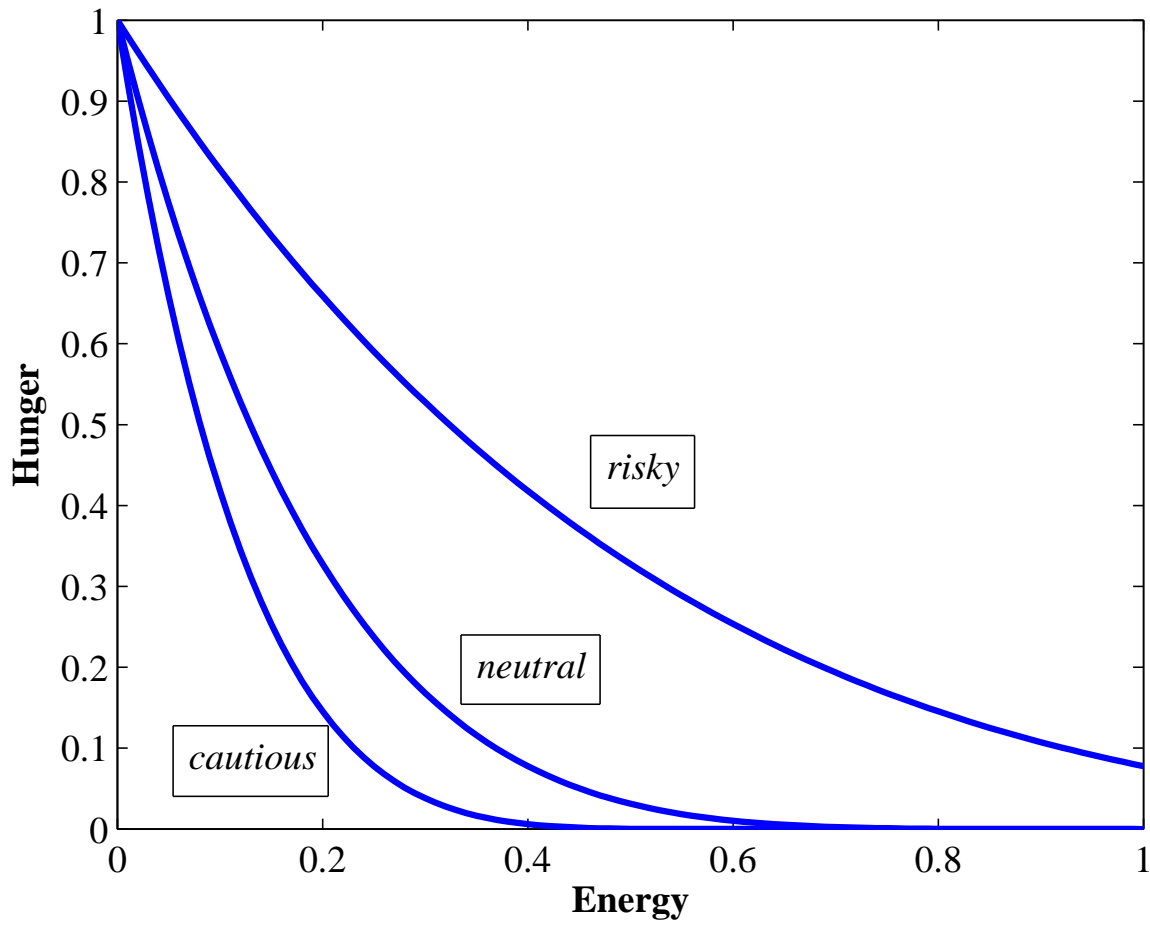


Figure 5: Hunger functions (Study 2).


in these latter two feedback conditions, the untrained beans were evaluated somewhat negatively (and even more so than in Study 1). However, a very different pattern emerges in the risky condition. Here the asymmetry between the learning of good and bad beans is effectively eliminated, and untrained beans are evaluated relatively positively. Analyses of variance (Feedback  $\times$  Hunger  $\times$  Valence for the number of correct choices and  $\Delta$ s; Feedback  $\times$  Hunger for the untrained patterns) indicated that all main effects and interactions were highly significant ( $p < .001$ ).

Table 2: Mean number of correct choices out of eighteen and mean absolute error ( $\Delta$ ) to good and bad input patterns, and mean evaluation of untrained patterns, in relation to hunger function and feedback (Study 2).


	N Correct Choices		Mean Absolute Error		Mean Evaluation
	Good	Bad	Good	Bad	Untrained
Full feedback	18.0	18.0	0.00	0.00	0.50
Contingent feedback	12.6	18.0	0.27	0.00	0.32
Confirmation bias	13.8	18.0	0.20	0.00	0.38
Cautious hunger					
Full feedback	18.0	18.0	0.00	0.00	0.51
Contingent feedback	12.0	18.0	0.30	0.00	0.27
Confirmation bias	12.1	18.0	0.30	0.00	0.28
Risky hunger					
Full feedback	18.0	18.0	0.00	0.00	0.50
Contingent feedback	18.0	18.0	0.00	0.00	0.51
Confirmation bias	18.0	17.4	0.00	0.01	0.54

These findings therefore demonstrate that the asymmetry in the learning of good and bad objects found in Study 1 could be replicated with a different network architecture. This asymmetry remained essentially unchanged when the hunger function was defined so as to have little or no effect except at low energy levels. However, where the action selection mechanism receives an extra boost of activation even when the network's energy level is high (risky condition), this appears sufficient to get the network to eat more beans provisionally categorized as bad, and so receive the feedback required to correct false negative beliefs.


These findings, however, disguise one important difference between the performances of the networks in the two studies. The data shown in Table 2 are based on 10 replications per cell at the end of 5000 epochs of training. However, the Study 2 network occasionally died early in training, i.e. reached zero




(a) Neutral hunger, contingent feedback




(b) Neutral hunger, confirmation bias




(c) Cautious hunger, contingent feedback



(d) Cautious hunger, confirmation bias



(e) Risky hunger, contingent feedback



(f) Risky hunger, confirmation bias

Figure 6: Mean evaluations for all attribute combinations as a function of hunger function and feedback.

energy before it had developed an adequate representation of the input space (essentially, so as to avoid eating too many bad beans). In the neutral condition, the numbers of deaths (i.e. extra runs required to produce 10 successful replications per cell) were 0, 0, and 12 respectively under full, contingent, and confirmation bias feedback, compared with 0, 0 and 8 in the cautious condition and 4, 9 and 23 in the risky condition. In short, the elimination of the learning asymmetry in the risky condition comes at the price of several deaths caused by indiscriminate eating. The combination of confirmation bias and the separation of the action-selection mechanism from the learning system architecture also seems to leave the network vulnerable to an early death. By contrast, in Study 1, there were no deaths in any of the conditions. A plausible explanation for the greater efficiency of the original network is that its energy unit operates almost as a part of the hidden layer with which it is connected, thus increasing the power of the hidden layer to discriminate between the different regions.

## 5 Discussion

Connectionist simulation has been employed effectively to generate theoretical insights in many areas of cognitive psychology. However, when extending this technique to social psychology, a major issue to be faced is that of how to represent the value individuals attach to particular objects. Network simulators are essentially programs for transforming particular abstract numerical patterns (vectors and matrices) into others. There is nothing intrinsically good or bad, or better or worse, about some numbers rather than others. We can, of course, choose to *define* (as here) activations of positive sign as standing for approval and activations of negative sign as standing for disapproval (e.g. Eiser et al., 1998). To do this, however, is just to adopt a mnemonic convention. There is nothing about these numbers as such that implies that anything *evaluative* or even *symbolic* is going on (and still less that computers can have attitudes).

Our approach, therefore, was not simply to show (as under full feedback) that connectionist networks can be trained to differentiate patterns that we have defined as standing for good and bad objects. Rather, we started by asking whether there may be anything in the process of learning itself that may distinguish how favorable and unfavorable attitudes are acquired. Our simulations were guided by the intuition that our attitudes are largely acquired by interaction with our environment and also that our attitudes guide such interaction. More specifically, we need to explore and approach objects in order to find out about them, but at the same time we are more likely to approach objects we expect to be good. Conversely, we will tend to avoid objects we expect to be bad unless motivated (here, by hunger) to engage in potentially risky exploration. Since, by avoiding such objects, we learn nothing that contradicts our initial aversion, our provisionally unfavorable attitude toward them will persist, and may even be strengthened. This corresponds to the classic finding in animal learning of avoidance behavior resisting extinction over time, and may account also for many human phobic behaviors and cognitions.

Our first aim was therefore to reproduce, in a highly restricted context, this asymmetry in the way we believe favorable and unfavorable attitudes are acquired. In Study 1, we modified the standard backpropagation of error algorithm by making updating of connection weights contingent on the network having produced an output above a specified threshold (equivalent to approach, or eating a bean). The effect of this was that, whereas the network still learnt the location of the bad beans, some good beans were never identified as such by the network. Modifying the algorithm further by including a confirmation bias for avoidance made this effect slightly stronger, and reduced the  $\Delta$  for the bad beans. The implication is that our ability to identify good objects may be incomplete, but that we are less likely to hold false positive than false negative beliefs. We also observed that the network in all conditions generalized its learning to new input patterns not previously presented during training. This is a consequence of it having acquired connection weights to solve the problem initially presented to it, and then applying these connection weights to new inputs. As can be seen from the landscape plots in Figures 3 and 6, untrained beans tended to be evaluated similarly to those in nearby regions that had been presented during training.

In Study 2, we employed a different computational architecture incorporating a distinction between a system for learning evaluations and a mechanism for selecting actions based on such evaluations. Evidently, the distinction between evaluation and behavior is fundamental to attitude theory. The network used in Study 2 makes this distinction more transparent. Variations between energy levels and response biases were also examined without assuming that hunger directly influenced the network's expectations regarding the valence of specific beans. Furthermore, the deterministic link between evaluation and action was modified by a stochastic probability function (noise). This meant that exploratory behavior could still occur occasionally even where beans were expected to be bad. Despite this important modification, the learning asymmetry observed in Study 1 was replicated (and, if anything strengthened) under two of the three hunger function conditions. This suggests that it may take more than an occasional contact with the truth to correct false negative beliefs. Part of the reason for this is that the network is not designed to learn the valence of each bean one at a time, but to form a distributed, i.e. configurational, representation of the input space as a whole. Put differently, the network associates valence with general *categories* of objects, and such categorical expectations appear robust enough to withstand occasional contradiction. However, when a risky hunger function was used (so that the network was still motivated to sample beans even when it had maximum energy), this extra boost toward exploratory or approach behavior was sufficient to eliminate the asymmetry between the learning of good and bad beans.

The effects of different hunger functions could be interpreted from the perspective of several theoretical approaches, including sensation-seeking (Zuckerman, 1994) and regulatory focus theory (Higgins, 1998). Our present research, however, does not attempt to simulate the processes underlying the development of individual differences in risk acceptance-aversion or approach-avoidance motivation (e.g. Elliot & Thrash, 2002). Rather, these effects evoke the classic distinction in signal detection theory (Swets, 1973) between sensitivity and response bias. Sensitivity refers to the ability of a system to discriminate reliably

between classes of objects, in this case, good and bad beans. Response bias refers to the tendency to set a criterion or response threshold at a level that involves acceptance of a higher number of either false positive or false negative errors, often so as to reflect the pay-off of potential benefits and costs. In Study 2, response bias was manipulated through predetermined hunger functions. However, our approach could be extended to consider how feedback from the environment might reinforce different exploration strategies (and hence lead to the acquisition of preferences for risk or caution), over and above its effects on evaluative expectancies that have been the focus of our present research. In any case, the findings from the connectionist modeling suggest that human performance in the learning situation would be improved by inducing participants to adopt a riskier approach to their exploratory behavior. That is, by more readily approaching beans about whose outcomes they are uncertain, participants should obtain more feedback and, hence, the learning asymmetry would be reduced. Effectively, such participants would be approximating a full feedback learning environment.

The aim of our simulation was to explore the implications of particular assumptions about the processes underlying the acquisition of attitudes. Our two studies demonstrate that asymmetries between positive and negative attitudes follow directly from relatively simple assumptions about the context in which people gain experience of their world. The most fundamental of these is that individuals make choices based on their expectations of outcomes, i.e. that they will approach things they expect to be good and give pleasure, and avoid things they expect to be bad and give pain. So long as these preconditions prevail, individuals who adopt an exploration strategy resembling that simulated here will tend to manage to identify sufficiently safe and rewarding regions of their life-space, albeit at the price of leaving some other potentially rewarding regions unexplored. Hence, for such individuals, positive experiences will tend to predominate over negative ones and, if we believe Parducci (1984), this will lead to feelings of happiness. Indeed, on average, people do seem to describe themselves as above average in happiness (Klar & Giladi, 1999) and positive traits (Hoorens, 1995), as well as less vulnerable to personal risks (Weinstein, 1989). Less encouragingly, though, people may persist in negative and prejudicial beliefs through a lack of any learning experience to contradict such beliefs. Even quite weak priming with negative beliefs can be self-reinforcing if individuals do not need to put the truth of their negative beliefs to the test. All this lends plausibility to the idea that we acquire attitudes, not so much to provide ourselves with a true and complete map of what is good and bad in our environment or life-space, but rather so that we can navigate through selected areas of that life-space with reasonable safety and gain. Attitudes, in short, are there to guide our behavior.

But of course, not all choices are that free. Even under ordinary circumstances, not all desirable outcomes can be attained and not all undesirable ones avoided. Very many individuals are subject to abusive and oppressive conditions where pain and punishment occur both frequently and inescapably. Research on learned helplessness (Abramson, Seligman & Teasdale, 1978) testifies to the damaging effects on individuals' well-being, motivation and self-esteem of uncontrollable negative events. While we are



proposing a view of attitude learning formulated at a high level of generality, we nonetheless readily acknowledge that there will be many contexts in which this assumption of free choice is less applicable. If individuals consistently fail to avoid negative events, positive experiences are unlikely to predominate over negative ones in their learning history, as Parducci (1984) assumes. If individuals lack the opportunity to achieve desired goals, they may persist in false positive beliefs that "the grass is greener" without ever being able directly to put these to the test. Such constraints could be modeled, but we have not done so here. The important point is that learning experiences of any kind can shape our evaluations of our environment and our own relation to it. Our simulations have focused on contexts where such learning experiences are themselves a function of evaluative beliefs.

## 6 References

- Abramson, L.Y., Seligman, M.E.P. & Teasdale, J.D. (1978). Learned helplessness in humans: Critique and reformulation. *Journal of Abnormal Psychology*, 87, 49-74.
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50, 179-211.
- Allport, G.W. (1935). Attitudes. In C. Murchison (Ed.), *Handbook of social psychology* (pp. 798-844). Worcester, MA: Clark University Press.
- Bargh, J.A. (1997). The automaticity of everyday life. In R.S. Wyer, Jr. (Ed.) *The automaticity of everyday life: Advances in Social Cognition*, Vol. X (pp. 1-62). Mahwah, NJ: Erlbaum.
- Baumeister, R.F., Bratslavsky, E., Finkenauer, C. & Vohs, K.D. (2001). Bad is stronger than good. *Review of General Psychology*, 5, 323-370.
- Edwards, W. (1954). The theory of decision making. *Psychological Bulletin*, 51, 380-417.
- Eiser, J.R. (1994). *Attitudes, chaos and the connectionist mind*. Oxford, Blackwell.
- Eiser, J.R., Claessen, M.J.A. & Loose, J.J. (1998). Attitudes, beliefs and other minds: Shared representations in self-organizing systems. In S.J. Read and L.C. Miller (Eds.) *Connectionist models of social reasoning and social behavior* (pp.313-354). Mahwah, NJ: Erlbaum.
- Elliot, A.J. & Thrash, T.M. (2002). Approach-avoidance motivation in personality: Approach and avoidance temperaments and goals. *Journal of Personality and Social Psychology*, 82, 804-818.
- Ellis, R. & Humphreys, G. (Eds.) (1999). *Connectionist psychology: A text with readings*. Hove: Psychology Press.
- Fazio, R.H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In M.P. Zanna (Ed.) *Advances in experimental social psychology* (Vol. 23, pp. 75-109). San Diego, CA: Academic Press.
- Fazio, R.H. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R.E. Petty and J.A. Krosnick, (Eds.). *Attitude strength: An-*

*tecedents and consequences* (pp. 247-282). Mahwah, NJ: Lawrence Erlbaum.

Fazio, R.H. (2001). On the automatic evaluation of associated evaluations: An overview. *Cognition and Emotion*, *15*, 115-141.

Fazio, R.H. & Eiser, J.R. (2000). *Attitude formation through associative learning: Valence asymmetries*. Paper presented to the annual meeting of the Society for Experimental Social Psychology, October. Atlanta, GA.

Fazio, R.H. & Zanna, M.P. (1981). Direct experience and attitude-behavior consistency. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 14, pp. 161-202). San Diego, CA: Academic Press.

Gurney, K. (1997). *An Introduction to Neural Networks*. London: UCL Press.

Harnad, S. (1987). (Ed.). *Categorical perception: The groundwork of cognition*. Cambridge: Cambridge University Press.

Heider, F. (1946). Attitudes and cognitive organization. *Journal of Psychology*, *21*, 107-112.

Hildum, D.C. & Brown, R.W. (1956). Verbal reinforcement and interviewer bias. *Journal of Abnormal and Social Psychology*, *53*, 108-111.

Higgins, E.T. (1998). Promotion and prevention: Regulatory focus as a motivational principle. In M.P. Zanna (Ed.) *Advances in Experimental Social Psychology*, Vol. 30 (pp. 1-46). San Diego, CA: Academic Press.

Hoorens, V. (1995). Self-favoring biases, self-presentation, and the self-other asymmetry in social comparison. *Journal of Personality*, *63*, 793-817.

Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science*, *79*, 2554-2558.

Insko, C.A. (1965). Verbal reinforcement of attitude. *Journal of Personality and Social Psychology*, *2*, 621-623.

Klar, Y. & Giladi, E.E. (1999). Are most people happier than their peers, or are they just happy? *Personality and Social Psychology Bulletin*, *25*, 585-594.

Latané, B. (1981). The psychology of social impact. *American Psychologist*, *36*, 343-356.

Lewin, K. (1936). *Principles of topological psychology*. New York: McGraw-Hill.

McClelland, J. & Rumelhart, D. (1988). *Explorations in parallel distributed processing*. Cambridge, MA: MIT Press.

McLeod, P., Plunkett, K. & Rolls, E.T. (1998). *Introduction to connectionist modelling of cognitive processes*. Oxford: Oxford University Press.

Nowak, A., Szamrej, J. & Latané, B. (1990). From private attitude to public opinion: A dynamic theory of social impact. *Psychological Review*, *97*, 362-376.

Nowak, A., Vallacher, R.R., Tesser, A. & Borkowski, W. (2000). Society of self: The emergence of collective properties in self-structure. *Psychological Review*, *107*, 39-61.

- Olson, M.A. & Fazio, R.H. (2001). Implicit attitude formation through classical conditioning. *Psychological Science*, *12*, 413-417.
- Page, M.M. (1974). Demand characteristics and the classical conditioning of attitudes experiment. *Journal of Personality and Social Psychology*, *30*, 468-476.
- Parducci, A. (1984). Value judgments: Toward a relational theory of happiness. In J.R. Eiser (Ed.) *Attitudinal judgment* (pp. 3-21). New York: Springer-Verlag.
- Petty, R.E. & Krosnick, J.A., Eds. (1995). *Attitude strength: Antecedents and consequences*. Mahwah, NJ: Lawrence Erlbaum.
- Read, S.J. & Miller, L.C. (1998). (Eds.) *Connectionist models of social reasoning and social behavior*. Mahwah, NJ: Erlbaum.
- Rozin, P. & Royzman, E.B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, *5*, 296-320.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), *Parallel Distributed Processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 318-362). Cambridge, MA: MIT Press.
- Seidenberg, M.S. (1993). Connectionist models and cognitive theory. *Psychological Science*, *4*, 228-235.
- Shultz, T.R. & Lepper, M.R. (1996). Cognitive dissonance as constraint satisfaction. *Psychological Review*, *103*, 219-240.
- Smith, E.R. (1996). What do connectionism and social psychology offer each other? *Journal of Personality and Social Psychology*, *70*, 893-912.
- Smith, E.R. & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, *4*, 108-131.
- Sutton, R.S. & Barto, A.G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: The MIT Press.
- Solomon, R.L. & Wynne, L.C. (1954). Traumatic avoidance learning: The principles of anxiety conservation and partial irreversibility. *Psychological Review*, *61*, 353-385.
- Staats, A.W. & Staats, C.K. (1958). Attitudes established by classical conditioning. *Journal of Abnormal and Social Psychology*, *57*, 37-40.
- Swets, J.A. (1973). The receiver operating characteristic in psychology. *Science*, *182*, 990-1000.
- Tolman, E.C. (1959). Principles of purposive behavior. In S. Koch (Ed.). *Psychology: A study of a science. Vol. 2. General systematic formulations, learning, and special processes*. New York: McGraw-Hill (pp. 92-157).
- Wegner, D.M. (1994). Ironic processes of mental control. *Psychological Review*, *101*, 34-52.

Weinstein, N.D. (1989). Optimistic biases about personal risks. *Science*, *246*, 1223-1232.

Zuckerman, M. (1994). *Behavioral expressions and biosocial bases of sensation seeking*. Cambridge: Cambridge University Press.