

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A Layered Network Model for Learning-to-learn and Configuration in Classical Conditioning

Permalink

<https://escholarship.org/uc/item/7815x1vw>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 8(0)

Author

Kehoe, E. James

Publication Date

1986

Peer reviewed

A Layered Network Model for Learning-to-Learn and Configuration
in
Classical Conditioning*

E. James Kehoe

School of Psychology, University of New South Wales

ABSTRACT

Networks composed of layers of adaptive elements provide a rigorous explanation for complex associative learning phenomena. In particular, a network composed of three adaptive elements can explain previously intractable phenomena, namely the rapid rate of reacquisitions, learning-to-learn, spontaneous configuration, and negative patterning (the exclusive-OR problem). This paper will compare the results of computer simulations to the behavioral results of classical conditioning experiments using the rabbit's nictitating membrane response.

INTRODUCTION

Layered networks of adaptive elements have featured prominently in contemporary theories of biological and machine cognition, particularly with regards to pattern recognition (Barto & Anderson, 1985; Feldman, 1985). Most notably, layered networks provide an elegant means for solving problems of nonlinear representation, for example, the exclusive-OR problem in which the system must learn to respond to each of two inputs but not their conjunction (Rumelhart, Hinton, & Williams, 1985). Rather than being "prewired" to represent particular combinations of inputs, layered adaptive networks of the appropriate type possess the ability to "tune" themselves to significant combinations of inputs (e.g., Barto, 1984; Barto, Anderson, & Sutton, 1982; Rumelhart et al., 1985). A less widely noted feature of layered networks is their ability to explain "learning-to-learn," that is a gain in the flexibility of the system's output as a consequence of prior training. At a more general level, a capacity for learning-to-learn may provide the

*This research was supported by Grant A28315236 from the Australian Research Grants Committee.

foundation for "insight" and other forms of "understanding" (Harlow, 1949). Accordingly, a major purpose of this paper is to demonstrate that some of the same features of layered networks that permit the recognition of arbitrary patterns also permit learning-to-learn.

This paper will present computer simulations of layered networks that are intended to duplicate the course of associative learning in a biological system, namely classical conditioning of the rabbit's nictitating membrane (NM) response (Gormezano, 1966; Gormezano, Kehoe, & Marshall, 1983). In the NM response preparation, the measured response is an extension of the third eyelid, which is innately elicited as an unconditioned response (UR) by a brief (50 ms) electrical-pulse unconditioned stimulus (US) administered to the surface of the skin posterior to the eye. Learning is produced by sequential presentations of a conditioned stimulus (CS) and the US, and, after a number of CS-US pairings, learning is evidenced by the acquisition of an NM conditioned response (CR) to the CS in advance of the US.

As a biological testbed for layered network models, classical conditioning procedures have several useful features:

- (1) Animals can be brought to the learning situation in a relatively naive state, thus approximating the initial state of an untutored network.
- (2) Animals do not require any prior verbal instructions, thus learning proceeds as a function of the stimulus inputs and response outputs that occur during the training task.
- (3) In classical conditioning procedures, it is possible to pose learning problems in a simplified way that can be duplicated in simple layered networks. For example, the exclusive-OR problem has its behavioral counterpart in Pavlov's (1927) negative patterning task. In that task, the animal is presented a mixture of three types of learning trials: (a) a tone CS that signals the US, (b) a light CS that also signals the US, and (c) a

compound tone + light stimulus that never signals the US. The animal can be said to have solved the negative patterning task when it generates CRs to the tone and the light but not the compound.

- (4) In many classical conditioning procedures, the CR appears to the CS in advance of the arrival of the US. By observing these anticipatory CRs, it is possible to trace the course of learning on a trial-by-trial basis. For purposes of testing a network model, the eventual achievement of a solution is perhaps less interesting than observing the intermediate states of the system prior to the solution state. For example, in solving the negative patterning problem, animals initially show considerable CR acquisition to the compound as well as the separate tone and light stimuli, after which responding to the compound gradually declines (e.g., Bellingham, Gillette-Bellingham, & Kehoe, 1985; Whitlow & Wagner, 1972; Woodbury, 1943).

In the remainder of this presentation, I shall describe in three stages a model of classical conditioning based on a layered network scheme. The model originates in those of Barto, Sutton, and their associates, which in turn are based on Hebb's (1949, 1972) theory of synaptic facilitation (e.g., Barto, 1984; Barto et al., 1982; Sutton & Barto, 1981). In brief, each stage of the model will encompass an increasing number of conditioning phenomena. The first-stage model will explain simple CR acquisition to one CS and a primitive form of learning-to-learn evidenced by progressive increases in the rate of successive acquisitions and extinctions conducted with the same CS (Hoehler, Kirschenbaum, & Leonard, 1973; Scavio & Thompson, 1979; Schmaltz & Theios, 1972; Smith & Gormezano, 1965). The second-stage model will encompass a more advanced form of learning-to-learn, namely a facilitation of CR acquisition to a new CS (e.g., light) after prior training with another, highly distinctive CS (e.g., tone) (Holt & Kehoe, 1985; Kehoe & Holt, 1984). Finally, the

third-stage model will explain a variety of simple pattern recognition phenomena, including negative patterning.

STAGE I: SUCCESSIVE ACQUISITIONS AND EXTINCTIONS

Figure 1 shows a schematic diagram of the network. The network contains two "sensory" elements, one for the tone CS (T) and one for the unconditioned stimulus (US). The output from T projects to an intermediate element (X), and the output from X projects to another element (R), which in turn gives rise to the observable behavior (CR/UR). Both nonsensory elements, namely X and R, receive an output from the US element.

Initial CR Acquisition

At the beginning of training with a naive animal, only the outputs from the US to X and R are effective. That is to say, only the US element can trigger an all-or-none firing of X and R. Initially, the T input is unable to trigger the intermediate element, but the T input does render its connection with X eligible for modification by the US input should it occur during a brief eligibility period that follows CS onset (Sutton & Barto, 1981). Thus, as the T-X connection strengthens over successive CS-US pairings, T will begin to trigger X. Then, the X-R connection will become eligible for change by the US's input to R. Observable CRs to the tone will only begin to appear when the intervening connections become strong enough so that T triggers X and then X triggers R. The changes in each of the interior connections, namely T-X and X-R, are governed by the linear operator process commonly used in current models of conditioning (Sutton & Barto, 1981). (See Appendix 1 for a full description of the implementation of the model). The firing of both X and R is all-or-none and is determined by a normally-distributed random threshold variable. Thus, on a given trial, X fires in response to an input from T only if the T-X connection weighting exceeds the threshold value on that trial. Likewise, R fires only if the X-R connection weight exceeds the current threshold.




Figure 1. A minimal network of two sensory elements (T, US) and two adaptive elements (X, R) for successive acquisitions and extinctions in classical conditioning.

Figure 2's left-hand column of panels shows the changes across blocks of CS-US trials in (a) the T-X connection, (b) the X-R connection, and (c) the percent CR measure produced by a computer simulation of the network's activities. As can be seen in the bottom panel, it is possible to reproduce a typical acquisition curve. The thresholds and growth rate parameters for both connections were selected so that the simulated curve would approximate the acquisition curve typically obtained in the rabbit NM response preparation under an 800-ms CS-US interval (see Appendix 1). As can be seen in the upper two panels, the T-X connection rises to a high level before the X-R connection shows any substantial change. For example, in the second block of training, the T-X connection was .69 while the X-R connection was only .10.

Subsequent Acquisitions

In its remaining panels, Figure 2 shows the simulated changes for the interior connections and percent CR across an initial extinction, a reacquisition, and a re-extinction. During




Figure 2. Simulation results for successive acquisitions and extinctions.

the initial extinction, the T-X connection declines at a steady rate, while the X-R connection declines to an asymptotic level of .70. As the T-X connection weakens and X's frequency of firing declines, the X-R connection is eligible for modification less and less often. In this way, the X-R connection is largely protected from extinction and thus remains intact. With respect to the simulated percent CR, it can be seen that the CR frequency reaches negligible levels while both the T-X and X-R connections are still appreciable. Consequently, during reacquisition in the third stage, both the T-X and X-R connections need relatively few reinforcements to rise to their asymptotic levels, yielding a relatively rapid rise in CR likelihood.

The second extinction does not appear particularly more rapid than the first extinction. To some extent, this simulated outcome is accurate; the available data suggest that the change in extinction rate is considerably slower than the change in acquisition rate across alternations of the training conditions.

Thus, this version of the model appears to be accurate to at least a first approximation.

STAGE II: LEARNING-TO-LEARN

Figure 3 shows an example of a learning-to-learn effect that has been repeatedly observed in conditioning of the rabbit NM response (Holt & Kehoe, 1985; Kehoe & Holt, 1984; Kehoe, Morrow, & Holt, 1984). In Phase I of this particular experiment, one group of rabbits received CS-US pairings in which the initial CS was an 800-ms tone. Another group served as a rest control. As




Figure 3. Example of initial CR acquisition to a tone CS (T+) and subsequent transfer to a light CS (L+). The point marked X indicates the initial response levels to the light CS.

expected, the former group showed CR acquisition to the tone, while the rest control group showed negligible responding. At the start of Phase II, both groups received four unreinforced presentations of an 800-ms light to determine the level of immediate cross-modal generalization. In the present experiment as in all our other studies, immediate transfer was not detectable. In Figure 3, the mean response likelihood on the tests are shown above the "X" marker on the abscissa. Only one animal, which happened to be in the pretrained group, responded twice to the light. However, once CS-US training with light was begun, the pretrained group showed extremely rapid CR acquisition to the new CS. For example, the animals in the pretrained group achieved a mean CR likelihood of 36% within the first block of reinforced light trials. By way of comparison, the naive animals in the control group achieved a mean CR likelihood of only 2% within the first block of reinforced light trials. In the present case, the pretrained group showed a higher level of responding to the light than the control group throughout Phase II. However, the asymptotic level of responding in the control group usually converges with that of the pretrained group. In other studies, we have shown that the positive transfer between tone and light is symmetric.

In order to explain the cross-modal learning-to-learn effect, it is only necessary to add an additional sensory element for the light to the network, as can be seen Figure 4. The input from light (L) projects to the intermediate element X just as the input from tone (T) does. Nothing else about the model is changed in any way.

Figure 5 shows the results of computer simulations for the learning-to-learn effect. The simulation of initial CR acquisition with the tone proceeds in the normal way for the model. In particular, observable CRs to the initial CS (tone) wait upon the successive strengthening of the T-X and X-R connections. However, in subsequent reinforced training with the new CS (light), the appearance of CRs requires only the establishment of the L-X connection, because the X-R connection has been already fully strengthened. Thus, as soon as the L-X

KEHOE




Figure 4. A minimal network of three sensory elements (T, L, US) and two adaptive elements (X, R) for learning-to-learn.




Figure 5. Simulation results for learning-to-learn.

KEHOE

connection becomes strong enough to trigger X, any firings of X triggered by L are immediately translated into observable CRs via the previously-established X-R connection. The rapid CR acquisition to the light is displayed in the learning curve for Phase II labelled as PRE, which denotes pretraining. In the way of contrast, a learning curve for a naive control condition is also displayed, labelled as REST. Thus, by relying on a common connection and the combination of convergent CS inputs, a layered network can explain the learning-to-learn effect.

In addition to demonstrating the learning-to-learn effect, my associates and I have found that it survives extinction of the original conditioned reflex (Kehoe et al., 1984). Figure 6 shows the results of one of our experiments. The key experimental group (4-E) received initial training at a 400-ms CS-US interval with one CS (CS1-US). Between CS1-US training and transfer training with a second CS (CS2-US), the animals in Group 4-E received a CS1-alone extinction procedure. Another experimental group (4-H) remained in their home-cages during the extinction procedure. In addition, two corresponding control groups (28-E and 28-H) initially received exposure to CS1 and the US but at a long 2,800-ms CS-US interval. Examination of the left-hand panel of Figure 6 reveals that Groups 4-E and 4-H showed conventional CR acquisition, while Groups 28-E and 28-H showed negligible levels of responding. The middle-panel shows that the Group 4-E displayed considerable extinction of the CR to CS1, whereas Group 28-E continued to display little responding. Finally, the right-hand panel reveals that, despite the near elimination of the initial conditioned reflex (CS1-CR) in Group 4-E, those animals acquired the new conditioned reflex (CS2-CR) as rapidly as their counterparts in Group 4-H, both of which showed positive transfer relative to their respective control groups.

On the theoretical side, the computer simulations successfully reproduced the ability of the learning-to-learn effect to survive disruption of the initial conditioned reflex. According to the computer simulations, the learning-to-learn effect survives for the same reasons that reacquisition after extinction is more rapid than initial acquisition. Figure 7




Figure 6. Learning-to-learn in Group 4-E survived extinction of the initial CR (Kehoe, Morrow, & Holt, 1984).




Figure 7. Simulation results for acquisition, extinction, and transfer training.

shows the results of computer simulations for the case in which there is an intervening extinction of the original conditioned reflex. As shown in Figure 7, the X-R connection is largely intact at the end of tone extinction. With the X-R connection still in place, pairings of the alternate CS (L) with the US can take advantage of the X-R connection and rapidly produce CRs as the L-X connection begins to strengthen. The lower right-hand panel of Figure 7 shows two simulated percent CR curves. The solid line represents acquisition to the light in the group that received tone pretraining followed by tone extinction (i.e., Group 4-E). The dotted line represents the simulated acquisition curve from a pretrained group that did not undergo extinction of the original conditioned reflex (i.e., Group 4-H). In agreement with the behavioral data, the two curves overlap perfectly.

STAGE III: CONFIGURAL LEARNING

The rabbit NM response preparation has expressed its sensitivity to patterns of multiple sensory inputs in a variety of ways. Figure 8 shows the course of differentiation between a compound and its components under three different training regimes. The lower panel shows the learning curves obtained under a negative patterning schedule, which corresponds to the exclusive-OR problem. As can be seen, differentiation proceeded slowly; responding to the compound, which was never followed by the US, declined only after extensive training (Bellingham et al., 1985). The upper right-hand panel reveals that differentiation proceeded much more rapidly in a positive patterning procedure, in which reinforced presentations of a tone + light compound (TL+) were intermixed with unreinforced presentations of the separate components (T-, L-) (Bellingham et al., 1985; Kehoe & Schreurs, in press). In logical terms, the positive patterning schedule corresponds to an AND problem. Differentiation of a compound from its components is not confined to procedures entailing explicit discrimination training. As shown in the upper left-hand panel, pairings of a compound with the US can produce spontaneous differentiation of the compound

KEHOE




Figure 8. Examples of CR acquisition to a tone, light, and compound (tone + light) in compound conditioning, positive patterning, and negative patterning procedures.

from its components (Kehoe, 1986; Kehoe & Schreurs, in press). On the basis of both explicit and implicit differentiation between a compound and its components, numerous theorists have proposed that the nervous system establishes distinctive encodings for the compound and its components, each with its own

excitatory or inhibitory associative strength (e.g., Bellingham et al., 1985; Hull, 1943, 1945; Kehoe & Gormezano, 1980; Razran, 1965, 1971; Rescorla, 1972, 1973; Whitlow & Wagner, 1972). While negative patterning clearly represents a nonlinear combination of the components, positive patterning and spontaneous differentiation may represent cases in which the CR-evoking capacity of the compound results from a linear combination of the CR-evoking capacities of the separate components, tone and light. Nevertheless, a history of reinforced exposure to a compound stimulus engages a combination process, linear or otherwise, that permits the subject to respond to the compound as a functional unit distinct from its components.

Figure 9 shows a schematic diagram of a network that can explain the configural learning phenomena. The network is essentially two parallel instances of the network used in the Stage II model. That is to say, the sensory inputs for tone, light, and the US project to a second intermediate element (Y), which in turn projects to the R element. For purposes of triggering an element by a joint input, it was assumed that the sum of currently eligible connection weights is compared to the element's threshold value. For changing the input weights, the present model followed the lead of Sutton and Barto (1981). In




Figure 9. A network of three sensory elements (T, L, US) and three adaptive elements (X, Y, R) for configural learning.

brief, when two inputs to either the X, Y, or R elements were simultaneously eligible for modification, then the inputs competed for the available connection weights supported by the US input. Thus, in training with a single CS, say the tone, the inputs from X and Y to the R element would compete with each other. In compound training with the tone and light, the inputs from T and L to the X element would compete with each other. Likewise, the T and L inputs to the Y element would compete with each other.

This competitive process was originally formulated to account for stimulus selection phenomena, in which increases in the weight of one stimulus input would be either blocked by prior increases in the weight of another concurrent input or overshadowed by more rapid increases in the weight of another concurrent input (e.g., Rescorla & Wagner, 1972; Sutton & Barto, 1981). However, this competitive process can also cause elements to become tuned to the combined T and L inputs. Specifically, competition between the T and L inputs would ensure that neither input by itself would gain sufficient connective weight to be able to reliably trigger the next element.

In order to discover a set of parameters that would accurately simulate configural learning, I manipulated two groups of parameters, namely the mean threshold value of each element (T_j) and the learning rate parameter for each element (a_j). Figure 10 shows the results of using the Stage III network to simulate the results of the compound stimulus experiments. The curves were obtained when (1) the X element had a higher learning rate than that of the Y element ($a_x = .100$, $a_y = .001$) and (2) the X element had a higher mean threshold than the Y element ($T_x = .65$, $T_y = .25$).

Inspection of Figure 10 reveals that the proposed model was able to simulate (a) the virtually complete differentiation between the compound and its components in compound conditioning, and (b) the relatively slow acquisition of negative patterning characterized by an initial rise in responding to the unreinforced compound stimulus followed by a gradual decline. The only detail missing in the simulations was the initial rise




Figure 10. Simulations of compound conditioning, positive patterning, and negative patterning.

in responding to the unreinforced components in the positive patterning schedule. However, in the rabbit NM response preparation, such a rise does not appear to be a universal feature of positive patterning (Bellingham et al., 1985). While the Stage III network was constructed to simulate the quantitative outcomes of the compound conditioning and patterning schedules, further simulation runs using the same parameter values have indicated that the Stage III network retains the basic characteristics and limitations of the competitive models

(Rescorla & Wagner, 1972; Sutton & Barto, 1981). In particular, the Stage III network reproduces blocking and conditioned inhibition, while being unable to generate latent inhibition.

DISCUSSION

The present simulations reveal that layered network models have considerable scope for explaining a variety of distinctive learning phenomena that have previously proved intractable to rigorous, elegant explanation. In particular, the network model not only reached the same endpoints as the behavioral phenomena but followed much the same course of acquisition. Although the particular network model used in this presentation was intended to be as nonspecialized as possible, it is nevertheless only an example of a larger class of models. For purposes of explaining rapid reacquisition and learning-to-learn phenomena, a large variety of layered network structures would be suitable. In an earlier version of the Stage II model, there were direct connections from each sensory input to the R element as well as to the intermediate element X. Simulations of that model revealed that it too could generate rapid reacquisition and learning-to-learn. These same phenomena should also appear under a huge range of rules for the adaptive elements, provided that the interior connections (e.g., X-R) do not change considerably faster than the connections from the sensory elements to the interior elements (e.g., T-X). However, successful simulation of the configural learning phenomena may occur only under a narrow range of network structures and adaptive rules, because these phenomena require a more delicate balancing of acquisition rates and threshold values to yield the appropriate connection weights.

REFERENCES

- Barto, A. G. (Ed.) (1984). Simulation experiments with goal-seeking adaptive elements. (AFWAL-TR-84-1022). Wright-Patterson AFB, OH: Avionics Laboratory, Air Force Wright Aeronautical Laboratories.
- Barto, A. G., & Anderson, C. W. (1985). Structural learning in connectionist systems. Paper presented to the Seventh Cognitive Science Conference.
- Barto, A. G., Anderson, C. W., & Sutton, R. S. (1982). Synthesis of nonlinear control surfaces by a layered associative search network. Biological Cybernetics, 43, 175-185.
- Bellingham, W. P., Gillette-Bellingham, K., & Kehoe, E. J. (1985). Summation and configuration in patterning schedules with the rat and rabbit. Animal Learning and Behavior, 13, 152-164.
- Feldman, J. A. (Ed.) (1985). Special issue on connectionist models and their applications. Cognitive Science, 9.
- Gormezano, I. (1966). Classical conditioning. In J. B. Sidowski (Ed.), Experimental methods and instrumentation in psychology (pp. 385-420). New York: McGraw-Hill.
- Gormezano, I., Kehoe, E. J., & Marshall, B. S. (1983). Twenty years of classical conditioning research with the rabbit. In J. M. Sprague and A. N. Epstein (Eds.), Progress in psychobiology and physiological psychology: Vol. 10 (pp. 197-275). New York: Academic Press.
- Harlow, H. F. (1949). The formation of learning sets. Psychological Review, 56, 51-65.
- Hebb, D. O. (1949). The organization of behavior: A neuropsychological theory. New York: John Wiley.
- Hebb, D. O. (1972). Textbook of psychology (3rd Edition). Philadelphia: W. B. Saunders.
- Hoehler, F. K., Kirschenbaum, D. S., & Leonard, D. W. (1973). The effects of overtraining and successive extinctions upon nictitating membrane conditioning in the rabbit. Learning and Motivation, 4, 91-101.
- Holt, P. E., & Kehoe, E. J. (1985). Cross-modal transfer as a function of similarities between training tasks in classical conditioning of the rabbit. Animal Learning and Behavior, 13, 51-59.

- Hull, C. L. (1943). Principles of behavior. New York: Appleton-Century-Crofts.
- Hull, C. L. (1945). The discrimination of stimulus configurations and the hypothesis of neural afferent interaction. Psychological Review, 52, 133-139.
- Kehoe, E. J. (1986). Summation and configuration in conditioning of the rabbit's nictitating membrane response to compound stimuli. Journal of Experimental Psychology: Animal Behavior Processes, 12, 186-195.
- Kehoe, E. J., & Gormezano, I. (1980). Configuration and combination laws in conditioning with compound stimuli. Psychological Bulletin, 87, 351-378.
- Kehoe, E. J., & Holt, P. E. (1984). Transfer across CS-US intervals and sensory modalities in classical conditioning of the rabbit. Animal Learning and Behavior, 12, 122-128.
- Kehoe, E. J., Morrow, L. D., & Holt, P. E. (1984). General transfer across sensory modalities survives reductions in the original conditioned reflex in the rabbit. Animal Learning and Behavior, 12, 129-136.
- Kehoe, E. J., & Schreurs, B. G. (in press). Compound-component differentiation as a function of CS-US interval and CS duration in the rabbit's conditioned nictitating membrane response. Animal Learning and Behavior.
- Pavlov, I. P. (1927). Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex (G. V. Anrep, trans.). London: Oxford University Press.
- Razran, G. (1965). Empirical codifications and specific theoretical implications of compound-stimulus conditioning: Perception. In W. F. Prokasy (Ed.), Classical conditioning (pp. 226-248). New York: Appleton-Century-Crofts.
- Razran, G. (1971). Mind in evolution. New York: Appleton-Century-Crofts.
- Rescorla, R. A. (1972). "Configural" conditioning in discrete-trial bar pressing. Journal of Comparative and Physiological Psychology, 79, 307-317.
- Rescorla, R. A. (1973). Evidence for the "unique stimulus" account of configural conditioning. Journal of Comparative and Physiological Psychology, 85, 331-338.

KEHOE

- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), Classical conditioning II (pp. 64-99). New York: Appleton-Century-Crofts.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1985). Learning internal representations by error propagation (ICS Tech. Rep. No. 8506). San Diego: University of California. Institute for Cognitive Science.
- Scavio, M. J., Jr., & Thompson, R. F. (1979). Extinction and reacquisition performance alternations of the conditioned nictitating membrane response. Bulletin of the Psychonomic Society, 13, 57-60.
- Schmaltz, L. W., & Theios, J. (1972). Acquisition and extinction of the classically conditioned response in hippocamectomized rabbits (Oryctolagus cuniculus). Journal of Comparative and Physiological Psychology, 79, 328-333.
- Smith, M., & Gormezano, I. (1965). Effects of alternating classical conditioning and extinction sessions on the conditioned nictitating membrane response of the rabbit. Psychonomic Science, 3, 91-92.
- Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. Psychological Review, 88, 135-171.
- Whitlow, J. W., & Wagner, A. R. (1972). Negative patterning in classical conditioning: Summation of response tendencies to isolable and configural components. Psychonomic Science, 27, 299-301.
- Woodbury, C. B. (1943). The learning of stimulus patterns by dogs. Journal of Comparative Psychology, 35, 29-40.

APPENDIX 1

SIMULATION OF THE LAYERED NETWORK MODELS

The computer simulation of the Stage I, II, and III networks was implemented in the following fashion:

1. The output of each element (FIRE_j) was either 1 or 0. The outputs of sensory elements T, L, and US were specified in the program on a trial by trial basis, while the outputs of X, Y, and R were determined by learning and output rules.

KEHOE

2. Each point of connection between successive elements had a connection weight designated as V_{ij} . The connection weights were designated as V_{tx} , V_{ty} , V_{lx} , V_{ly} , V_{xr} , and V_{yr} , where the first letter in the subscript refers to the origin of the output and the second letter refers to the element receiving the output. These connection weights started at zero and were unbounded in both the positive and negative directions. The connection weights from the US output to the X, Y, and R elements were fixed at 1.00.

3. Each trial was divided into two time steps, the CS period and the US period.

4. During the CS period, the following events occurred:

a. The output of the T and L elements was determined by the programmed trial sequence, and the appropriate connections with X and Y became eligible for change. For example, on a compound trial, there were outputs from T and L. Accordingly, the T-X, T-Y, L-X, and L-Y connections all became eligible for change.

b. A separate threshold (T_j) was independently determined for X, Y, and R by generating a random number between 0.00 and 0.99. Across trials, the distribution of thresholds was approximately normal. To alter the threshold, a constant was added or subtracted from the random number.

c. The output of X, Y, and R was:

$$\text{FIRE}_j = 1 \text{ if sum of eligible input weights} > T_j$$

$$\text{FIRE}_j = 0 \text{ otherwise.}$$

For example, on a compound trial, the output of X was determined by comparing $V_{tx} + V_{lx}$ against T_x . At the same time, the output of Y was determined by comparing $V_{ty} + V_{ly}$ against T_y .

d. After the outputs of X and Y were determined, then the output of R was determined by same process. For example, if only X fired on a particular compound trial, then the output of R (FIRE_r) was determined by comparing V_{xr} against T_r .

5. During the US period, the eligible connection weights throughout the network were modified according to the

KEHOE

Rescorla-Wagner model, where the change in a connection weight (dV_{ij}) followed the formula:

$$dV_{ij} = a_j (L_j - \sum V_{ij}),$$

where

a_j was the rate of change parameter for the target element of the connection. ($0 < a_j < 1$). On non-reinforced trials, a_j was modified by parameter B_0 ($0 < B_0 < 1$) (See Point 7 below).

L_j was the total connection strength of eligible connections on a target element that could be supported by the US input on any given trial ($L_j = 1.0$ on reinforced trials. $L_j = 0.0$ on nonreinforced trials.)

$\sum V_{ij}$ was the net sum of the associative strengths of all concurrent eligible inputs to the j th element.

6. In order to compute the CR likelihood for the tone CS, light CS, and compound CS for each block of training trials, a series of 30 "phantom CS periods" was conducted at the end of each block. Thus, the CS period for each type of trial was repeatedly conducted and the likelihood of a CR was determined. However, the change formula applicable during the US period of training trials was not used. Effectively, these phantom CS periods served as repeated test trials but without the extinctive effect that a prolonged series of test trials would have had in a behavioral experiment.

7. The simulations of successive acquisitions, extinctions, and learning-to-learn (Figures 2, 5, 7) used the following parameter values: $a_x = a_r = 0.02$, mean $T_x = 0.75$, mean $T_r = 0.50$, $B_0 = 0.15$. The simulations for compound conditioning, positive patterning, and negative patterning (Figure 10) used the following parameter values: $a_x = 0.100$, $a_y = 0.001$, $a_r = 0.004$, mean $T_x = 0.65$, mean $T_y = 0.25$, mean $T_r = 0.50$, $B_0 = 0.33$.