

Brief Papers

Equilibria of Perceptrons for Simple Contingency Problems

Michael R. W. Dawson and Brian Dupuis

Abstract—The contingency between cues and outcomes is fundamentally important to theories of causal reasoning and to theories of associative learning. Researchers have computed the equilibria of Rescorla–Wagner models for a variety of contingency problems, and have used these equilibria to identify situations in which the Rescorla–Wagner model is consistent, or inconsistent, with normative models of contingency. Mathematical analyses that directly compare artificial neural networks to contingency theory have not been performed, because of the assumed equivalence between the Rescorla–Wagner learning rule and the delta rule training of artificial neural networks. However, recent results indicate that this equivalence is not as straightforward as typically assumed, suggesting a strong need for mathematical accounts of how networks deal with contingency problems. One such analysis is presented here, where it is proven that the structure of the equilibrium for a simple network trained on a basic contingency problem is quite different from the structure of the equilibrium for a Rescorla–Wagner model faced with the same problem. However, these structural differences lead to functionally equivalent behavior. The implications of this result for the relationships between associative learning, contingency theory, and connectionism are discussed.

Index Terms—Artificial neural networks, associative learning, contingency.

I. INTRODUCTION

A fundamental characteristic of an adaptive agent is the ability to detect causal relations [1]. However, the real world poses constant challenges to this ability, because cues do not signal outcomes with complete certainty [2]. It has been argued that adaptive systems deal with worldly uncertainty, whether these systems are humans [3] or animals [4], [5], by becoming “intuitive statisticians.” The notion of “intuitive statistician” has been rigorously developed in a series of important papers to mean sensitivity to contingency, where contingency is defined in a normative model as a contrast between conditional probabilities [1], [6]–[11]. For instance, consider the simple situation that is detailed in the contingency table provided in Table I. The contingency between the cue and the outcome is formally defined as the difference in conditional probabilities ΔP , where $\Delta P = P(O|C) - P(O|\sim C)$ [6].

Manuscript received October 12, 2011; revised January 12, 2012 and April 10, 2012; accepted April 20, 2012. The work of M. R. W. Dawson was supported in part by an NSERC Discovery Grant.

The authors are with the Department of Psychology, University of Alberta, Edmonton, AB T6G 2P9, Canada (e-mail: mdawson@ualberta.ca; bdupuis@ualberta.ca).

Digital Object Identifier 10.1109/TNNLS.2012.2199766

More sophisticated models, such as the probabilistic contrast model (e.g., [8]) or the power PC theory [7], define more complex probabilistic contrasts that are possible when multiple cues occur and when what they signal depends upon the context in which they are considered.

Because many associative learning paradigms can be interpreted as teaching contingencies to humans or animals, another issue that has arisen in the literature is the relationship between formal contingency theories and formal theories of associative learning [12]. In particular, researchers have compared the predictions of the Rescorla–Wagner model of learning [13] to formal theories of contingency [1], [7], [14]. This is typically accomplished by determining the equilibria for the Rescorla–Wagner model, and then comparing the associative strengths of the Rescorla–Wagner model at equilibrium with probabilistic contrasts defined by contingency theory. An equilibrium of the Rescorla–Wagner model is a set of associative strengths defined by the model at the point where changes in error defined by Rescorla–Wagner learning asymptote to zero [15]. While in some instances the Rescorla–Wagner model predicts the conditional contrasts defined by a formal contingency theory like the power PC model, in other situations it fails to generate these predictions [9].

The formal results relating contingency theory to the Rescorla–Wagner model have been assumed to also apply to connectionist models of associative learning [1], [5]. Researchers have claimed that there is a formal equivalence [16]–[18] between learning as defined by the Rescorla–Wagner model and learning as defined by the so-called delta rule, which is an error-correcting method that is used to train simple artificial neural networks [19], [20]. Such claims are used to support the informal conclusion that any results pertaining to the relationship between the Rescorla–Wagner model and contingency theory also apply to artificial neural networks trained with the delta rule. That is, if for at least some cases x , the Rescorla–Wagner model and contingency theory are equivalent, and if the Rescorla–Wagner model is equivalent to delta rule learning, then it seems safe to conclude that for these same cases x , networks trained with the delta rule should be equivalent to contingency theory.

One example of this indirect argument is provided by Cheng [9], who performs a detailed computational analysis of the relationship between the Rescorla–Wagner model and contingency theory. She emphasizes the Rescorla–Wagner model because “the learning rule it incorporates is a version of the “delta rule” commonly used in connectionist models. My analysis of this model should therefore be relevant to connectionist models using this rule, whatever the content domain of the model” [9, p. 371]. However, Cheng neglects to conduct a computational analysis that directly relates contingency theory to artificial neural networks. Cheng and

TABLE I

SIMPLE CONTINGENCY SITUATION IN WHICH A CUE CAN OCCUR (C) OR NOT ($\sim C$), AND AN OUTCOME CAN OCCUR (O) OR NOT ($\sim O$) AS WELL

	O	$\sim O$
C	a	b
$\sim C$	c	d

Note: The four letters in the table represent the frequency of co-occurrence of the two types of events. Using this table, $\Delta P = P(O|C) - P(O|\sim C) = a/(a+b) - c/(c+d)$.

Holyoak [1] provide a second example of the indirect argument when they use the assumed equivalence between the delta rule and the Rescorla–Wagner model to define asymptotic associative weights for the latter. A third example of this indirect argument is provided by Shanks’ [5] interpretation of the computation work of Chapman and Robbins [14]. In an appendix to their article, Chapman and Robbins prove that in a particular situation “the Rescorla–Wagner model reduces to the ΔP rule” [14, p. 545]. However, Shanks [5, p. 112] uses the indirect argument to interpret the proof in a connectionist light, claiming that “Chapman and Robbins have established the very important fact that the delta rule, at asymptote, yields weights that are identical to ΔP .”

However, there are important reasons to be wary of using the assumed relationship between the Rescorla–Wagner model and the simple artificial neural networks to infer relationships between networks and contingency theory. First, previous proofs of the formal equivalence between the Rescorla–Wagner model and the delta rule [16]–[18] neglect to include a critical component of artificial neural networks trained by the delta rule—the nonlinear activation function that converts an output unit’s net input into activation. Dawson [23] has shown that these proofs assume a linear relationship between net input and activity, and therefore do not apply to simple neural networks such as the traditional perceptron [21], [22] that uses a step function to compute output unit activity, or a modern variation of the perceptron [19], [23] that uses a logistic activation function. When the nonlinear activation function is taken into account, a formal relationship between the Rescorla–Wagner model and the complete (i.e., nonlinear) networks can still be established [23]. However, the inclusion of the activation function imposes a crucial algorithmic difference between network learning and the Rescorla–Wagner model: the activation function serves as a theory of how internal associations are converted into network behavior, while a theory of behavior is not part of the Rescorla–Wagner formalism [13], [24]. As a result, a perceptron that uses the logistic activation can generate different behaviors than can a model trained using the Rescorla–Wagner model, and in many cases can overcome some limitations faced by the Rescorla–Wagner model [23].

In short, the relationship between Rescorla–Wagner learning and artificial neural network learning is more complicated than one might expect from older comparisons [16], [18]. As a result, it is unwise to use these older analyses as the basis for an indirect link between networks and contingency theory. Instead, computational analyses that directly explore

the relationships between connectionist networks and contingency theory are required. The purpose of this brief is to provide one such analysis. It is proven below that when a simple artificial neural network reaches equilibrium for a basic contingency theory problem, this equilibrium appears to be quite different from the equilibrium of the Rescorla–Wagner model for the same contingency problem. That is, in contrast to Shanks’ [5] connectionist interpretation of Chapman and Robbins’ [14] proof, the connection weights of the network are not identical to ΔP . However, ΔP can be recovered by comparing the behavior of the network in different cue situations.

II. DERIVING THE EQUILIBRIUM

A. Derivation

To begin, consider the simple contingency problem that was presented earlier in Table I. Chapman and Robbins [14] proved that when Rescorla–Wagner learning reaches equilibrium for this problem the associative strength between the cue and the outcome was exactly equal to ΔP . Their proof required the assumption that there were two cues involved, the one of interest (C) that was present on some trials and absent on others (as in Table I), and a second (X) that represented cues from an experimental context that were present on every trial. Rescorla–Wagner learning would alter the strengths of two associations, the one between C and the outcome (V_C) and the one between X and the outcome (V_X). For the situation defined in Table I, Chapman and Robbins found that at equilibrium $V_C = a/(a+b) - c/(c+d)$. Let us now proceed to derive the equilibrium for a perceptron faced with the same contingencies.

One can train a simple perceptron on the Table I contingency problem. The perceptron would have a single input unit that would be turned on with a value of 1 when C is present, and turned off with a value of 0 when C is absent. This signal would be sent through a single connection, with connection weight w_c , to a single output unit. The desired response of this output unit would be 1 in trials in which the outcome O occurred, and would be 0 in trials in which O did not occur. On any given trial, the net input *net*, (i.e., the total signal) to the output unit is equal to w_c times the activation value of the input unit. A nonlinear transformation of the net input produces the output unit’s response for the trial. Let us define this nonlinear transformation as the logistic equation, which is an activation function that is commonly employed in artificial neural networks [19], [25], [26]

$$f(\text{net}) = \frac{1}{1 + e^{-(\text{net}) - \theta}}. \quad (1)$$

The logistic equation is often described as a “squashing” function, because it is a sigmoid-shaped function that squashes values of net input, which can range from negative to positive infinity, into the range from 0 to 1. In (1), *net* is the net input from the perceptron’s input unit, and θ is a constant that is called the bias of the logistic equation. When net input equals θ , the logistic equation returns a value of 0.5. θ can be described as the value of a weight between an “extra” input unit and the output unit, where this “extra” input unit has

an activation value of 1 for every pattern that the perceptron is presented. In other words, the use of θ in the logistic equation is equivalent to Chapman and Robbins' [14] use of an extra cue to represent the constant presence of experimental context [23].

Assume that when the cue is present, the logistic activation function computes an activation value that we will designate as o_c , and that when the cue is absent it returns the activation value designated as $o_{\sim c}$. We can now define the total error of responding for the perceptron [i.e., its total error for the $(a + b + c + d)$ number of patterns that represent a single "sweep" in which each instance of the contingency problem given in Table I is presented once]. For instance, on a trial in which C and O both occur (i.e., both C and O equal 1), the perceptron's error for that trial is the squared difference between O and o_c . As there are a of these trials, the total contribution of this type of trial to overall error is $a(1 - o_c)^2$. Applying this logic to the other three cells of Table I, overall error E can be defined as follows:

$$E = a(1 - o_c)^2 + b(0 - o_c)^2 + c(1 - o_{\sim c})^2 + d(1 - o_{\sim c})^2 \\ = a(1 - o_c)^2 + b(o_c)^2 + c(1 - o_{\sim c})^2 + d(1 - o_{\sim c})^2. \quad (2)$$

For a perceptron to be at equilibrium, it must have reached a state in which the error term defined in (2) has been optimized, so that error can no longer be decreased by using the learning rule to alter the perceptron's weight. To determine the equilibrium of the perceptron for the Table I problem, we begin by taking the derivative of (2) with respect to the activity of the perceptron when the cue is present (o_c). This derivative is presented as (3). We also need to determine the derivative of (2) with respect to the activity of the perceptron when the cue is not present ($o_{\sim c}$). This derivative is presented as (4)

$$\frac{\partial E}{\partial o_c} = 2(a(o_c - 1) + bo_c) \quad (3)$$

$$\frac{\partial E}{\partial o_{\sim c}} = 2(C(o_{\sim c} - 1) + do_{\sim c}). \quad (4)$$

One condition of the perceptron at equilibrium is that o_c is a value that causes the derivative in (3) to be equal to 0. In (5), this derivative is set to 0 and the equation is solved to determine the value of o_c . The reader will note that this value is equal to $a/(a + b)$, which is equal to the conditional probability $P(O|C)$

$$0 = 2(a(o_c - 1) + bo_c) \\ = a(o_c - 1) + bo_c \\ = ao_c - a + bo_c \\ a = o_c(a + b) \\ \frac{a}{a + b} = o_c \\ P(O|C) = o_c. \quad (5)$$

A second condition of the perceptron at equilibrium is that $o_{\sim c}$ is a value that causes the derivative in (4) to be equal to 0. In (6), this derivative is set to 0 and the equation is solved to determine the value of $o_{\sim c}$. The reader will note that this value is equal to $c/(c + d)$, which is equal to the

conditional probability $P(O|\sim C)$

$$0 = 2(c(o_{\sim c} - 1) + do_{\sim c}) \\ = c(o_{\sim c} - 1) + do_{\sim c} \\ = co_{\sim c} - c + do_{\sim c} \\ c = o_{\sim c}(c + d) \\ \frac{c}{c + d} = o_{\sim c} \\ P(O|\sim C) = o_{\sim c}. \quad (6)$$

To provide a concrete example of the implications of these equations, let us consider the result of training a perceptron on a "toy problem" consistent with Table I. Imagine a training set consisting of 20 patterns, each involving a single cue represented by the activation of a perceptron that has only one input unit. The cue is present in exactly half of these patterns, and is reinforced (i.e., the perceptron is trained to output a value of 1.0) for eight of these training patterns, and is not reinforced (i.e., the perceptron is trained to output a value of 0.0) for the remaining two patterns. The cue is absent in the remaining 10 patterns, 2 of which are reinforced, while the remaining 8 are not reinforced. This statement of the problem permits the four entries of Table I to be filled out as follows: $a = 8$, $b = 2$, $c = 2$, and $d = 8$. For these table values, $\Delta P = (a/(a + b)) - (c/(c + d)) = (8/(8 + 2)) - (2/(2 + 8)) = 0.6$. Using software developed in our laboratory [27], a gradient descent rule was used to train a perceptron on this problem using a learning rate of 0.1, with the bias of the output unit and the connection weight randomly initiated in the range $[-0.1, 0.1]$. Four hundred training epochs, in which each of the 20 patterns is presented once in random order, were conducted; after 400 epochs, the network had stabilized. At the end of this training, the weight of the connection between the input unit and the output unit was 2.76, and the bias of the output unit was -1.38 . When the cue was presented by turning the input unit on, an output value of 0.8 was generated, which is $P(O|C)$. When the cue was not presented by turning the input unit off, an output value of 0.2 was presented, which is $P(O|\sim C)$.

B. Implications

One implication of the proof developed above is that for the type of contingency problem described in Table I, at equilibrium, the output of a perceptron trained on this problem can literally be described as a conditional probability. When the cue is present, perceptron output can be literally interpreted as the likelihood of the outcome given the cue. Similarly, when the cue is absent, perceptron output can be literally interpreted as the likelihood of the outcome in the absence of the cue. This was shown in the toy example provided above, where the perceptron activity was equal to the appropriate conditional probability depending upon the presence or absence of the cue.

This result makes contact with the extensive empirical literature on probability matching. Probability matching occurs when the probability with which an agent makes a choice among alternatives mirrors the probability associated with the outcome or reward of that choice [28]. Studies involving a

variety of subjects, including insects, fish, turtles, pigeons, and humans have not only shown the existence of probability matching but have also shown that probability matching is adaptive: when the probability of reinforcement associated with a cue changes, the choice probabilities exhibited by the agent are quickly adjusted [29]–[36]. It was recently shown that perceptrons that use the logistic activation function match probabilities, and also quickly adapt these probabilities when reinforcement contingencies are altered [37]. The proof above grounds this empirical finding in mathematics by demonstrating that perceptron outputs are identical to conditional probabilities.

A second implication of the proof developed above is that an equilibrium for a perceptron faced with the Table I contingency problem is not, as expected by Shanks [5], identical to the equilibrium for the Rescorla–Wagner model. At equilibrium, the associative strength for the cue C that is determined by Rescorla–Wagner training is literally ΔP . This is not the case for the perceptron. This was shown, for instance, in the example given above; in the network that was trained, neither the connection weight nor the bias was equal to ΔP .

The fact that the associative strengths at equilibrium for the Rescorla–Wagner model differ from those at equilibrium for the perceptron does not indicate qualitative differences between the two in the context of the contingency problem being solved. That is, the two systems achieve equilibria that appear to be different because the two systems use associative strengths in different ways to produce behavior (i.e., to generate judgments of contingency). For the Rescorla–Wagner model, the general assumption is that associative strengths are converted into responses by a linear transformation [23]. Thus, if the behavior of such a model is to reflect ΔP , then ΔP must be directly represented in associative strengths, as proved by Chapman and Robbins [14]. In contrast, the perceptron uses a nonlinear transformation when it converts associative strengths into responses. Therefore ΔP cannot be directly encoded as a connection weight. Instead, ΔP must be computed after a response is generated—by taking the difference between a perceptron’s output when the cue is present and the same when the cue is absent. For instance, in the example provided earlier, if after training, one takes the difference between perceptron activity when the cue is present (0.8) and perceptron activity when the cue is absent (0.2), the result is 0.6, which is the value of ΔP given the representation of that problem in Table I format.

It might be argued that a proper difference between the two equilibria has not been established because one is framed in terms of associative strength, while the other is framed in terms of perceptron output. However, the value of θ and the value of the connection weight w_c can easily be computed given the results in (5) and (6). First, if one sets the value of the logistic function in (1) to $c/(c+d)$, assumes $net = 0$, and solves for θ , then it is found that θ equals $\ln(d/c)$. Second, if one sets the value in (1) to $a/(a+b)$, assumes $\theta = \ln(d/c)$, and solves for w_c , then it is found that w_c equals $\ln(d/c) - \ln(b/a)$ —which is not equal to ΔP . (One can solve for w_c in this case because in this simple network, when $C = 1$, $net = w_c$.)

A third implication of the proof developed above is that one cannot naively assume that the formal equivalence of Rescorla–Wagner learning and delta rule learning [16]–[18] also establishes that the Rescorla–Wagner model is identical to a connectionist network like the perceptron. The analysis of the perceptron’s equilibrium reveals a final state that is structurally quite different from that predicted from Shanks’ [5] interpretation of the Chapman and Robbins’ [14] proof. That is, for the perceptron, ΔP is not directly represented as a connection weight.

This simply suggests that further formal research is required to directly establish the relationship between contingency theory and artificial neural networks. Modern contingency theory is concerned with contrasts between probabilities in situations involving multiple cues, and Danks [15] has demonstrated how equilibria for Rescorla–Wagner models can be computed in multiple-cue situations. Future formal research is required to determine equilibria for artificial neural networks in multiple-cue situations in order to investigate the degree of agreement or disagreement between networks and contingency theory. Beginning such work with the study of simple perceptrons is likely to bear fruit, because these simple networks are still the source of surprising and interesting results [38], [39], and because the behavior of perceptrons in multiple-cue situations suggests that this simple kind of network can mimic core empirical regularities. For instance, one key aspect of adaptive animal behavior is using multiple cues to maximize survival, and to use changes in the information provided by multiple cues to modify behavior accordingly [4]. Perceptrons have been shown to demonstrate such abilities, for instance, by reacting to new combinations of multiple cues to modify response probabilities in a navigation task [40]. It would be expected that formal analyses of the equilibria of such networks would shed a great deal of insight about their relation to more sophisticated versions of contingency theory.

A fourth implication of our results follows from the third: if naïve assumptions about the equivalence between the Rescorla–Wagner and neural networks are incorrect (as we have demonstrated), then a more rigorous account of the relationship is likely to shed new insights into the relationships between Rescorla–Wagner learning, neural network models, and contingency theory. In particular, mathematical knowledge concerning neural networks may provide new approaches to understanding learning about contingency.

For example, the proof developed above was based on a quadratic definition of network error, because this formulation of error has been central to studying the relation between Rescorla–Wagner and neural network learning [16], [18], [23]. However, other definitions of error are possible [39]. For instance, some researchers have suggested that network error for noisy or stochastic environments might be better characterized in terms of measures of entropy [41]–[44], or equivalently using error metrics that maximize information [45]–[47]. Future research that explores the relationships between contingency theory, animal learning, and neural networks using the mathematics of information theory is likely to produce interesting and important results.

REFERENCES

- [1] P. W. Cheng and K. J. Holyoak, "Complex adaptive systems as intuitive statisticians: Causality, contingency, and prediction," in *Comparative Approaches to Cognitive Science*, H. L. Roitblat and J.-A. Meyer, Eds. Cambridge, MA: MIT Press, 1995, pp. 271–302.
- [2] J. Dewey, *Experience and Nature*, 2nd ed. Chicago, IL: Open Court Publishing, 1929.
- [3] C. R. Peterson and L. R. Beach, "Man as an intuitive statistician," *Psychol. Bulletin*, vol. 68, no. 1, pp. 29–46, Jul. 1967.
- [4] C. R. Gallistel, *The Organization of Learning*. Cambridge, MA: MIT Press, 1990.
- [5] D. R. Shanks, *The Psychology of Associative Learning*. Cambridge, U.K.: Cambridge Univ. Press, 1995.
- [6] L. G. Allan, "A note on measurement of contingency between two binary variables in judgment tasks," *Bulletin Psychonomic Soc.*, vol. 15, no. 3, pp. 147–149, 1980.
- [7] P. W. Cheng, "From covariation to causation: A causal power theory," *Psychol. Rev.*, vol. 104, pp. 367–405, Apr. 1997.
- [8] P. W. Cheng and L. R. Novick, "A probabilistic contrast model of causal induction," *J. Personal. Social Psychol.*, vol. 58, pp. 545–567, Apr. 1990.
- [9] P. W. Cheng and L. R. Novick, "Covariation in natural causal induction," *Psychol. Rev.*, vol. 99, pp. 365–382, Apr. 1992.
- [10] R. A. Rescorla, "Pavlovian conditioning and its proper control procedures," *Psychol. Rev.*, vol. 74, no. 1, pp. 71–80, 1967.
- [11] R. A. Rescorla, "Probability of shock in the presence and absence of CS in fear conditioning," *J. Comparat. Physiol. Psychol.*, vol. 66, no. 1, pp. 1–5, 1968.
- [12] D. R. Shanks, "Associationism and cognition: Human contingency learning at 25," *Quart. J. Experim. Psychol.*, vol. 60, pp. 291–309, Mar. 2007.
- [13] R. A. Rescorla and A. R. Wagner, "A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement," in *Classical Conditioning II: Current Research and Theory*, A. H. Black and W. F. Prokasy, Eds. New York: Appleton-Century-Crofts, 1972, pp. 64–99.
- [14] G. B. Chapman and S. J. Robbins, "Cue interaction in human contingency judgment," *Memory Cognit.*, vol. 18, pp. 537–545, Sep. 1990.
- [15] D. Danks, "Equilibria of the Rescorla–Wagner model," *J. Math. Psychol.*, vol. 47, pp. 109–121, Apr. 2003.
- [16] M. A. Gluck and G. H. Bower, "From conditioning to category learning - an adaptive network model," *J. Experim. Psychol.-General*, vol. 117, pp. 227–247, Sep. 1988.
- [17] M. A. Gluck and C. Myers, *Gateway to Memory: An Introduction to Neural Network Modeling of the Hippocampus and Learning*. Cambridge, MA: MIT Press, 2001.
- [18] R. S. Sutton and A. G. Barto, "Toward a modern theory of adaptive networks: Expectation and prediction," *Psychol. Rev.*, vol. 88, no. 2, pp. 135–170, 1981.
- [19] M. R. W. Dawson, *Minds and Machines: Connectionism and Psychological Modeling*. Malden, MA: Blackwell, 2004.
- [20] G. O. Stone, "An analysis of the delta rule and the learning of statistical associations," in *Parallel Distributed Processing*, vol. 1, D. E. Rumelhart and J. McClelland, Eds. Cambridge, MA: MIT Press, 1986, pp. 444–459.
- [21] F. Rosenblatt, "The perceptron: A probabilistic model for information storage and organization in the brain," *Psychol. Rev.*, vol. 65, no. 6, pp. 386–408, 1958.
- [22] F. Rosenblatt, *Principles of Neurodynamics*. Washington, DC: Spartan Books, 1962.
- [23] M. R. W. Dawson, "Connectionism and classical conditioning," *Comparat. Cognit. Behavior Rev.*, vol. 3, pp. 1–115, 2008, doi: 10.3819/ccbr.2008.30008.
- [24] R. R. Miller, "Challenges facing contemporary associative approaches to acquired behavior," *Comparat. Cognit. Behavior Rev.*, vol. 1, pp. 77–93, Jan. 2006.
- [25] W. Bechtel and A. A. Abrahamsen, *Connectionism and the Mind: Parallel Processing, Dynamics, and Evolution in Networks*, 2nd ed. Malden, MA: Blackwell, 2002.
- [26] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, pp. 533–536, Oct. 1986.
- [27] M. R. W. Dawson, *Connectionism: A Hands-On Approach*, 1st ed. Malden, MA: Blackwell, 2005.
- [28] N. Vulkan, "An economist's perspective on probability matching," *J. Econ. Surv.*, vol. 14, no. 1, pp. 101–118, Feb. 2000.
- [29] E. R. Behrend and M. E. Bitterman, "Probability-matching in the fish," *Amer. J. Psychol.*, vol. 74, no. 4, pp. 542–551, Dec. 1961.
- [30] W. K. Estes and J. H. Straughan, "Analysis of a verbal conditioning situation in terms of statistical learning theory," *J. Experim. Psychol.*, vol. 47, no. 4, pp. 225–234, Apr. 1954.
- [31] M. E. Fischer, P. A. Couvillion, and M. E. Bitterman, "Choice in honeybees as a function of the probability of reward," *Animal Learn. Behavior*, vol. 21, pp. 187–195, Aug. 1993.
- [32] V. Graf, D. H. Bullock, and M. E. Bitterman, "Further experiments on probability-matching in the pigeon," *J. Experim. Anal. Behavior*, vol. 7, no. 2, pp. 151–157, 1964.
- [33] T. Keasar, E. Rashkovich, D. Cohen, and A. Shmida, "Bees in two-armed bandit situations: Foraging choices and possible decision mechanisms," *Behavioral Ecol.*, vol. 13, pp. 757–765, Nov.–Dec. 2002.
- [34] K. L. Kirk and M. E. Bitterman, "Probability-learning by the turtle," *Science*, vol. 148, no. 3676, pp. 1484–1485, Jun. 1965.
- [35] N. Longo, "Probability-learning and habit-reversal in the cockroach," *Amer. J. Psychol.*, vol. 77, no. 1, pp. 29–41, 1964.
- [36] Y. Niv, D. Joel, I. Meilijson, and E. Ruppin, "Evolution of reinforcement learning in uncertain environments: A simple explanation for complex foraging behaviors," *Adapt. Behavior*, vol. 10, no. 1, pp. 5–24, 2002.
- [37] M. R. W. Dawson, B. Dupuis, M. L. Spetch, and D. M. Kelly, "Simple artificial neural networks that match probability and exploit and explore when confronting a multiarmed bandit," *IEEE Trans. Neural Netw.*, vol. 20, no. 8, pp. 1368–1371, Aug. 2009.
- [38] M. Fernandez-Delgado, J. Ribeiro, E. Cernadas, and S. B. Ameneiro, "Direct parallel perceptrons (DPPs): Fast analytical calculation of the parallel perceptrons weights with margin control for classification tasks," *IEEE Trans. Neural Netw.*, vol. 22, no. 11, pp. 1837–1848, Nov. 2011.
- [39] S. Raudys, R. Kybartas, and E. K. Zavadskas, "Multicategory nets of single-layer perceptrons: Complexity and sample-size issues," *IEEE Trans. Neural Netw.*, vol. 21, no. 5, pp. 784–795, May 2010.
- [40] M. R. W. Dawson, D. M. Kelly, M. L. Spetch, and B. Dupuis, "Using perceptrons to explore the reorientation task," *Cognition*, vol. 114, no. 2, pp. 207–226, 2010.
- [41] E. B. Baum and F. Wilczek, "Supervised learning of probability distributions by neural networks," in *Neural Information Processing Systems*, D. Z. Anderson, Ed. New York: AIP, 1988, pp. 52–61.
- [42] J. J. Hopfield, "Learning algorithms and probability distributions in feedforward and feedback networks," *Proc. Nat. Acad. Sci. United States Amer.*, vol. 84, pp. 8429–8433, Dec. 1987.
- [43] S. A. Solla, E. Levin, and M. Fleisher, "Accelerated learning in layered neural networks," *Complex Syst.*, vol. 2, no. 6, pp. 625–639, Dec. 1988.
- [44] B. S. Wittner and J. S. Denker, "Strategies for teaching layered networks classification tasks," in *Neural Information Processing Systems*, D. Z. Anderson, Ed. New York: AIP, 1988, pp. 850–859.
- [45] M. Plumbley, "Information-theoretic approaches to neural network learning," in *Neural Network Perspectives on Cognition and Adaptive Robotics*, A. Browne, Ed. Bristol, U.K.: Institute Physics Publishing, 1997, pp. 72–90.
- [46] M. D. Plumbley, "Information processing in negative feedback neural networks," *Netw.-Comput. Neural Syst.*, vol. 7, pp. 301–305, May 1996.
- [47] M. D. Plumbley, "Do cortical maps adapt to optimize information density?" *Netw.-Comput. Neural Syst.*, vol. 10, pp. 41–58, Feb. 1999.