Challenges from Probabilistic Learning for Models of Brain and Behavior

Nicolás Marchant, Enrique Canessa, and Sergio E. Chaigneau

Abstract Probabilistic learning is a research program that aims to understand how animals and humans learn and adapt their behavior in situations where the pairing between cues and outcomes is not always completely reliable. This chapter provides an overview of the challenges of probabilistic learning for models of the brain and behavior. We discuss the historical background of probabilistic learning, its theoretical foundations, and its applications in various fields such as psychology, neuroscience, and artificial intelligence. We also review some key findings from experimental studies on probabilistic learning, including the role of feedback, attention, memory, and decision-making processes. Finally, we highlight some of the current debates and future directions in this field.

Keywords Probabilistic learning · Category learning · Feedback · Decision-making · Cognitive models

1 Introduction

For a very long time, behavioral scientists have been wondering how animals can learn and adapt their behavior to environmental demands. One big concern among some scientists was that the pairing between cues and outcomes is not always completely reliable. If we honor Darwin's hypothesis, animals would seek to adapt their behavior to the environmental demands, even in unreliable situations. Since the early days of behaviorism, this research program concerned with how animals

N. Marchant (✉) · S. E. Chaigneau
Center for Social and Cognitive Neuroscience, School of Psychology, Universidad Adolfo Ibáñez, Santiago, Chile
e-mail: nicolas.marchant@edu.uai.cl; sergio.chaigneau@uai.cl

E. Canessa
Faculty of Engineering and Science, Universidad Adolfo Ibáñez, Viña del Mar, Chile
e-mail: ecanessa@uai.cl

T. Veloz et al. (eds.), *Trends and Challenges in Cognitive Modeling*, STEAM-H: Science, Technology, Engineering, Agriculture, Mathematics & Health,
https://doi.org/10.1007/978-3-031-41862-4_6

can learn under unreliable outcomes conditions was known as probabilistic learning (Brunswik 1943; Edwards 1961; Castellan 1973).

To study probabilistic learning is quite simple. The experimenter has only to adjust some schedule routine in order to make the pairing between cues and outcomes unreliable. For example, one of the first conditioning experiments that relied on the use of probabilistic learning was the experiment of Rescorla (1968). He adjusted the pairing routines between a tone (conditioned stimulus; CS) and a shock (unconditioned stimulus; US) by implementing different conditions in which the US was probabilistically paired with the CS. He noted that the conditioning strength was higher whenever the US and the CS occurred deterministically (i.e., with 100% of matching accuracy), but strength continually declined as the pairing became more unreliable (e.g., 80% of accurate matchings; 60% of accurate matchings). Overall, probabilistic associations between US and CS impoverish the conditioned learning. However, simple conditioning experiments were limited in their capacity to explain complex human behaviors such as inference, reasoning, and categorization. So, in the cognitive turn (see Miller 2003 for a brief review), conditioning was abandoned and replaced by cognitive explanations of probabilistic learning (Estes 1976; Lindell 1976). The cognitive turn on probabilistic learning highlighted the relationship between learning and memory formation, and whether those learned associations can be recovered through explicit processing.

A way of formalizing how memory and learning occur in the cognitive system is by means of computational modeling. Since the beginning of the cognitive turn, it is common to find mathematical formulations that represent a certain cognitive function as a testable "algorithmic" hypothesis (see Wilson and Collins 2019). This means that experimental and computational modeling explanations of probabilistic learning work together. In the current work, we will be interested in those computational/mathematical explanations. However, our discussion will be kept at a general level, so that readers that are not familiar with this area of research can grasp the general problems of the field. For those readers interested in the gory details, we provide numerous references.

We recognize that there are different areas of cognition interested in probabilistic processing and how to model it, such as decision-making (Tversky and Kahneman 1974), reasoning (Oaksford and Chater 2007), or Bayesian learning (Tenenbaum et al. 2006). However, our primary focus here is on the process of learning to classify when receiving probabilistic feedback across a sequence of consecutive trials. By understanding human probabilistic learning, it should be possible to better develop artificial learning systems that face the same problems as humans do (e.g., such as in machine learning; Frénay and Verleysen 2014). Also, it should be possible to develop a better explanation of the environmental demands that people face in natural learning conditions (e.g., doctors learning to diagnose from a set of symptoms; Estes 1986). This chapter presents some of the most well-known challenges that research on probabilistic learning faces in the era of cognitive neuroscience.

2 Challenges to Probabilistic Learning Modeling

2.1 Error Correction

One of the first challenges faced by modelers and experimentalists in explaining probabilistic learning was whether people actually learn from their errors and what is the mechanism that people and animals use to correct their behavior. As we outlined in the introduction, discoveries from behaviorisms stated that animals learn because of the teaching signal, which can be either an aversive or appetitive stimuli. One of the first mathematical formulations of error correction was the Rescorla and Wagner (1972) learning rule. In short, this learning rule states that the associative strength of a cue, which is stored in memory, depends on its predictive value of the outcome (López and Shanks 2008). Subjects will assign specific associative strengths depending on the cue predictiveness, which is acquired through learning, and thus, a cue with a highly predictive value will gain more associative strength. A way in which this model is often discussed is by saying that learning will be enhanced if the outcome is surprising, i.e., it is not predicted by the currently active cues. This surprisingness is computed through an error term which represents the discrepancy between the expectation and the actual status of the outcome in the current trial. The Rescorla and Wagner model (1972) was of special interest because it addresses problems that earlier theories of conditioning seemed unable to explain. One of those problems is the blocking effect (Kamin 1969), which states that when two cues are presented within the same trial, subjects will only learn the cue with the most predictive value while blocking the cue with the lower predictive value. Also, the model is consistent with neuroanatomical evidence regarding how the dopamine system works in humans and animals. Research carried out by Schultz (1999) shows evidence that dopaminergic neurons in the macaque midbrain fire when the predictive cue is presented (which was previously paired with a response), predicting the same firing rate that will trigger the response. Since then, these results have been replicated many times, with human and nonhuman animals (Daw and Doya 2006). Because the model provides an account for experiments with different salient cues and also because of its biological plausibility, the Rescorla and Wagner model is included, in one way or another, in many learning models.

A probabilistic learning model that incorporates the Rescorla and Wagner rule was the configural cue model developed by Gluck and Bower (1988a, b), Gluck (1991). They implemented a connectionist model (see Thomas and McClelland 2008) to address whether people can learn two potential categories from a set of different cues. This kind of experimental procedure is known as Multiple Cue Probability Learning (MCPL; Edgell 1980; Estes 1986). Here, the categories (i.e., Disease R and Disease C) are the outcomes which provide probabilistic feedback to different combinations of four different cues (i.e., the symptoms). Across 250 trials (of different combinations of symptoms), there was a probability of 25% that a specific trial will be correctly classified as Disease R (and a 75% of correctly be classified as Disease C). However, the first symptom (i.e., the bloody nose) had a 69% diagnosticity for Disease R (and the opposite for Disease C). The authors found that subjects learned that the first symptom was indeed more diagnostic of Disease R, independent of the base-rate of Disease R being low (i.e., 25%). Later, the authors modeled the behavioral results using their configural cue model. Using an error correction activation algorithm (the least mean squared rule (LMS), which as Sutton and Barto (1981) noted, is a special case of the Rescorla and Wagner rule), their model successfully predicted that the first symptom should become, in fact, more diagnostic. This occurs because different cues and the combination of those cues compete with each other to match the teaching signal, while the winning cue is the one that reduces the error between the response and the outcome (i.e., the classification).

However, the configural cue model was incapable of explaining all classification phenomena. Criticisms from Nosofsky and colleagues argued that the model was insufficient to account for rule-based classification (Nosofsky et al. 1994). This led to the emergence of other computational models which integrate an error term that were able to explain some rule-based problems. One such model is the ALCOVE (Kruschke 1992) model, which is an exemplar-based model embedded with the same LMS algorithm function. In brief, ALCOVE, much like other exemplar models, assumes that different stimuli are stored in memory as individual traces. Computationally, this idea is implemented as a neural net where a hidden layer has nodes representing each of the exemplars in the training set. However, ALCOVE still retains a learning mechanism which updates attentional resources by trial and error. Exemplar-based models were preferred, because they integrate a cognitive explanation of the attentional resources combined with an error-correction mechanism.

2.2 Feedback Discounting

Despite preferences in the literature for exemplar-based models of learning and categorization, they faced a second major challenge. A phenomenon known as feedback discounting (or error discounting) pushed the limits of these kinds of models. Feedback discounting occurs because people (and perhaps even nonhuman animals) will eventually accept a certain level of unavoidable error, and, continually, they will begin to discount feedback information slowing down their learning (Estes 1984; Kruschke and Johansen 1999; Craig et al. 2011).

Kruschke and Johansen (1999) developed an extension of the ALCOVE model which is able to capture whenever people stop using feedback information. They called this new model RASHNL (Rapid Attention SHifts 'N' Learning). This model incorporates a feedback discounting mechanism whenever attention is shifted away from irrelevant cues (those that produce error) to relevant cues (those that reduce error). Thus, in a trial-by-trial fashion, the RASHNL model will eventually reduce the learning rate from which it updates the cue-outcome associations (note that the RASHNL is an extension of the ALCOVE model, using the same error-correction term to update learning). Furthermore, the feedback discounting phenomenon has received supporting evidence from electrophysiology studies (EEG). A study conducted by Sewell and colleagues showed that participants who discount feedback entirely eliminate an EEG component known as feedback-related negativity (FRN), while participants who do not discount feedback presented a standard FRN frequency (Sewell et al. 2018). The FRN component usually elicits a peak evoked signal between 200 and 300 ms after the presentation of the feedback, being generally larger for negative feedback rather than for positive feedback (Cohen et al. 2011). In a nutshell, probabilistic learning models should take feedback discounting into account. As noted by Craig et al. (2011), models of probabilistic learning tend to improve whenever they incorporate a feedback discounting mechanism. Moreover, participants who discounted feedback showed different brain signals often found in the middle region of the EEG scalp. Regardless of the previous evidence, it is still unclear what conditions cause participants to start feedback discounting, and whether this is an automatic process or an explicit conscious strategy.

2.3 Normative Responses

Researchers in the psychology of decision-making and behavioral economics were interested in knowing whether people behaved according to normative criteria when dealing with probabilistic information. If so, then it should be found that people behave close to normatively. From the decision-making literature, it has been suggested that maximizing is the normative response when dealing with uncertainties (Fiorina 1971; Shanks et al. 2002). In short, maximizing states that people will always place a certain item into the response that is most likely to belong to. For example, if we have an item s, in which 80% of the trials belong to keyboard response A, then, people should maximize their responses by always responding A whenever item s is presented. However, there is a debate whether people always behave according to maximizing, which brings us closer to our third challenge. For some researchers, people often deviate from maximizing, and instead, they rely on a suboptimal response strategy which is called probability matching (Castellan 1973; Friedman and Massaro 1998; Shanks et al. 2002). Probability matching states that people will progressively match their responses according to the outcome criteria. For example, if item s belongs to response A with an 80% chance, then, subjects will tend to respond 80% of the times that item s belongs to response A.

There is still a debate under which circumstances people will respect maximizing or will fall into probability matching. For example, Shanks et al. (2002) created different experimental situations in which probability matching would be undesirable. However, there were always people (albeit a small proportion) that relied on probability matching. Studies that were concerned about how people learned to categorize under unreliable situations showed that – on average – people follow a probability matching pattern across learning trials (Little and Lewandowsky 2009a, b; Craig et al. 2011; Sewell et al. 2018). Obviously, this pattern of results places challenges for cognitive modeling, considering that maximizing and probability matching seem so different in terms of possible underlying cognitive mechanisms.

2.4 Cognitive Processing

So far, we have reviewed three challenges that are critical when researchers want to develop a probabilistic learning model. Whether our model has to update learned responses using an error correction algorithm, or whether it should be implemented with a feedback discounting mechanism whenever people stop relying on the informativeness of the feedback, or whether our model follows normative responses or deviates greatly from them. However, such challenges tell us little about the cognitive mechanisms underlying probabilistic learning. Error correction algorithms are thought to rely on associative-based processing, in which motor responses are guided through contingency routines of stimulus-outcome associations (Gluck and Bower 1988a, 1988b; Gluck 2008; Marchant et al. 2022; Marchant and Chaigneau 2022). On the other hand, it is not clear which cognitive processes underlie the phenomena of feedback discounting and probability matching. For some authors, they are explicit rules; for others, they are implicit rules, so the debate is still open. Our fourth challenge is related to what cognitive processes underlie probabilistic learning and how modeling might help us to understand such processes.

Evidence in the 1990s showed that amnesic patients perform similar to controls in a probabilistic learning task known as the Weather Prediction Task[1] (WPT), but just for the first 50 trials. Later in training, normal control samples outperform amnesic patients. Knowlton and colleagues believed that this effect occurred, because control subjects were capable of formulating a declarative strategy which they maintained "online" through the course of learning, while amnesic patients were incapable of doing that (Knowlton et al. 1994, 1996a; Meeter et al. 2006). However, a different set of evidence on Huntington disease patients, a disease that affects mostly the basal ganglia and other subcortical structures and reduces motor control and motor planning, showed that Huntington's disease patients also perform poorly in the WPT (Knowlton et al. 1996b). Thus, both, a motor-based component and a rule-based strategy, might be both necessary to learn under probabilistic feedback conditions.

Gluck et al. (2002) wondered whether there are different kinds of strategies that people rely on when solving probabilistic learning. Implementing a WPT with a debriefing phase just after the experiment ended, they asked subjects to verbally

[1] In the Weather Prediction Task, subjects are presented with combinations of playing cards with different patterns (i.e., geometric figures) combinations. The subjects must learn to use them to predict the weather (i.e., rain or sun). During training, subjects are presented with combinations of cards (i.e., one to four cards), while each specific combination is probabilistically associated with the two outcomes.

report how they solved the task. They found that most participants relied on a singleton strategy early in training (i.e., subject learned the optimal response for each of the four possible patterns of the WPT and guessed on the remaining trials). However, they also found that some participants changed their strategy in the later phases of learning toward a one-cue strategy (i.e., responding based on the presence or absence of one single cue) or to a multi-cue strategy (i.e., responding based on something like probability matching). Behavioral and patient-based evidence support that probabilistic learning relies on associative and motor processing and also on suboptimal strategies which can be retrieved through verbal reports. In the next and final challenge, we return to the idea of whether those cognitive processes occur simultaneously or compete with each other. Obviously, very different computational models ensue from each of these alternatives.

2.5 Rule-Based or Associative Mechanisms?

There is an ongoing debate in the literature about whether probabilistic learning is explained by a rule-based, associative-based, or a mixture between the two processes. Some researchers believe that rule-based processing comprises explicit declarative knowledge, whereas associative-based processing comprises implicit automatic learning (Ashby and O'Brien 2005). This distinction is often referred to as the dual-systems view in psychology, and certainly, it is not common only in probabilistic learning. There is evidence of dual-systems in reasoning (Sloman 1996) and also in decision-making (Evans 2008). In the field of probabilistic learning (also referred to as probabilistic categorization), computational models have been used to attempt to understand under what conditions subjects perform one or the other kind of processing.

A well-known computational model that embodies the dual-systems view is the COVIS model (Ashby et al. 1998, 2007; Ashby and Crossley 2012). COVIS assumes that these two systems (i.e., associative or procedural and declarative or rule-based) compete with each other to account for the best results according to task demands. One system is the procedural system in which there is little explicit verbal access or awareness of implicit memories. This system is feedback dependent and relies on the use of implicit memory systems (Maddox et al. 2004; Smith and Grossman 2008). The second system is a declarative system, which is engaged whenever a rule-based task is performed. This system maintains a series of subprocesses such as selecting, focusing, and switching rules. COVIS is not the only model that assumes a competition between two processing systems. Another learning model is ATRIUM (Attention to Rules and Instances in a Unified Model; Erickson and Kruschke 1998). This model assumes that people compute both rule-based strategies and exemplar-based computations (similar to ALCOVE and RASHNL models mentioned above). Similar to the COVIS model, Erickson and Kruschke (1998) conclude that people will rely on the use of rules or exemplar-based processing depending on task demands.

However, other researchers have been skeptical regarding the dual-view competing mechanism hypothesis (see Newell et al. 2011) and have questioned whether probabilistic learning involves some sort of self-awareness during the experimental task. If that is the case, then it is unlikely that the two processes occur independently, and some kind of interaction must be occurring (Evans et al. 2003). A study carried out by Lagnado et al. (2006) inspected if subjects rely on the use of self-insight (i.e., being able to report their own thought processes) under a probabilistic categorization problem. By implementing the WPT paradigm, the authors tracked learning across trials using a rolling regression method. This statistical method showed that regression coefficients accurately tracked the statistical contingencies over the task, revealing that subjects accurately learned the task structure. But, more importantly, the coefficient weights also correlated with self-insight regarding how much they used a certain cue. Thus, self-insight regression weights correlate with objective task performance, revealing a kind of interaction between an explicit processing (how useful a subject believes that a certain cue is) with an implicit process (how much a subject has learned about the cue-outcome contingencies).

Furthermore, other studies wondered if other cognitive processes typically involved in rule-based behavior would have an impact on probabilistic learning. A study carried out by Newell et al. (2007) wondered whether working memory capacity, which is usually related to explicit or rule-based behavior, will have an impact during a probabilistic learning task. They showed that a concurrent memory task will impoverish performance on the WPT, meaning that performance on the WPT is dependent on the use of working memory (i.e., on explicit rules). Another research by Rolison et al. (2011) showed that by using an MCPL (multiple cue probabilistic learning task) task, working memory capacity was only useful whenever a negative cue (i.e., a cue negatively correlated with the outcome) was presented. Whenever a positive cue was presented (i.e., a cue positively correlated with the outcome), working memory appeared to be unnecessary. Undoubtedly, learning models have helped us to understand the cognitive basis of probabilistic learning, making it possible to test different hypotheses. The debate is still open as to whether both mechanisms (associative and rule-based) compete with each other or whether there is a kind of interaction between the two. Future experiments and models should address under what circumstances one kind of mechanism interferes with the other.

3 Summary

Since the earlier days of behaviorism, the study of probabilistic learning has been a major endeavor, encompasing learning. And also to more complex experimental situations involving decisions that people often face in their daily lives. Cognitive explanations contributed to the development of formal computational models that represent how learning occurs in the mind and how it relates to memory. In this book chapter, we outlined some of the most discussed challenges faced by researchers when modeling probabilistic learning.

A first reviewed challenge was whether our model has to integrate an *error correction* algorithm that updates past responses according to a teaching signal. A typical error correction algorithm used in connectionist models is the LMS algorithm, which is closely related to Rescorla and Wagner (1972) learning rule. A second challenge was whether our model should incorporate a *feedback discounting* mechanism. Evidence has shown that people often stop considering feedback information in situations where it becomes too unreliable for learning. A third challenge is whether subjects respond *normatively* (i.e., maximizing) or whether they tend to use probability matching. There is a large discussion about whether people maximize their responses or whether they match the probability of the outcome. Evidence suggests that most subjects under most situations rely on probability matching. A fourth challenge is regarding the *cognitive processing* that underlies probabilistic learning. There is still a debate whether probabilistic learning is explained by an associative-based system or whether it is explained by the use of logical rules (declarative system). Some researchers have proposed that people use a range of possible different strategies according to individual parameters (i.e., focusing on one stimulus, on combination of stimuli, and so on). And a final reviewed challenge, closely related to the previous one, is whether or not *associative* and *rule-based* are *competing mechanisms* or not. Some authors have suggested that implicit and explicit processing compete in order to achieve a good performance, while other researchers have been skeptical to this idea proposing that probabilistic learning is explained by an interaction between declarative memory (self-insight) and implicit processing.

Certainly, the experiments and modeling of probabilistic learning have improved our understanding of the human mind. It addresses questions regarding how people interact, process, and store in memory the environment's unreliable information; what are the neural mechanisms that our brain uses when learning probabilistic information; and how we have to develop formal mathematical models that integrate algorithms that enable us to explain behavior.

References

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A Neuropsychological Theory of Multiple Systems in Category Learning. *Psychological Review*, *105*(3), 442–481. https://doi.org/10.1037/0033-295X.105.3.442

Ashby, F. G., & Crossley, M. J. (2012). Automaticity and multiple memory systems. *Wiley Interdisciplinary Reviews: Cognitive Science*, *3*(3), 363–376. https://doi.org/10.1002/wcs.1172

Ashby, F. G., Ennis, J. M., & Spiering, B. J. (2007). A Neurobiological Theory of Automaticity in Perceptual Categorization. *Psychological Review*, *114*(3), 632–656. https://doi.org/10.1037/0033-295X.114.3.632

Ashby, F. G., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Sciences*, *9*(2), 83–89. https://doi.org/10.1016/j.tics.2004.12.003

Brunswik, E. (1943). Organismic achievement and environmental probability. *Psychological review*, *50*(3), 255–272. https://doi.org/10.1037/h0060889

Castellan, N. J. (1973). Multiple-cue probability learning with irrelevant cues. *Organizational Behavior and Human Performance*, *9*(1), 16–29. https://doi.org/10.1016/0030-5073(73)90033-0

Craig, S., Lewandowsky, S., & Little, D. R. (2011). Error discounting in probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(3), 673–687. https://doi.org/10.1037/a0022473

Cohen, M. X., Wilmes, K., & van de Vijver, I. (2011). Cortical electrophysiological network dynamics of feedback learning. *Trends in Cognitive Sciences*, *15*(12), 558–566. https://doi.org/10.1016/j.tics.2011.10.004

Daw, N. D., & Doya, K. (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, 16(2), 199–204. doi:https://doi.org/10.1016/j.conb.2006.03.006

Edgell, S. E. (1980). Higher order configural information processing in nonmetric multiple-cue probability learning. *Organizational Behavior and Human Performance*, *25*(1), 1–14. https://doi.org/10.1016/0030-5073(80)90022-7

Edwards, W. (1961). Probability learning in 1000 trials. *Journal of Experimental Psychology*, *62*(4), 385–394. https://doi.org/10.1037/h0041970

Erickson, M. A., & Kruschke, J. K. (1998). Rules and Exemplars in Category Learning. *Journal of Experimental Psychology: General*, *127*(2), 107–140. https://doi.org/10.1037/0096-3445.127.2.107

Estes, W. K. (1976). The cognitive side of probability learning. *Psychological Review*, *83*(1), 37–64. https://doi.org/10.1037/0033-295X.83.1.37

Estes, W. K. (1984). Global and local control of choice behavior by cyclically varying outcome probabilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*(2), 258–270. https://doi.org/10.1037/0278-7393.10.2.258

Estes, W. K. (1986). Array models for category learning. *Cognitive Psychology*, *18*(4), 500–549. https://doi.org/10.1016/0010-0285(86)90008-3

Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, *59*, 255–278. https://doi.org/10.1146/annurev.psych.59.103006.093629

Evans, J. S. B., Clibbens, J., Cattani, A., Harris, A., & Dennis, I. (2003). Explicit and implicit processes in multicue judgment. *Memory & Cognition*, *31*(4), 608–618. https://doi.org/10.3758/BF03196101

Fiorina, M. P. (1971). A note on probability matching and rational choice. *Behavioral Science, 16*, 158–166.

Frénay, B., & Verleysen, M. (2014). Classification in the presence of label noise: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, *25*(5), 845–869. https://doi.org/10.1109/TNNLS.2013.2292894

Friedman, D., & Massaro, D. W. (1998). Understanding variability in binary and continuous choice. *Psychonomic Bulletin and Review*, *5*(3), 370–389. https://doi.org/10.3758/BF03208814

Gluck, M. A. (1991). Stimulus generalization and representation in adaptive network models of category learning. *Psychological Science*, 2(1), 50–55. https://doi.org/10.1111/j.1467-9280.1991.tb00096.x

Gluck, M. A. (2008). Behavioral and neural correlates of error correction in classical conditioning and human category learning. In Gluck, M. A, Anderson, J. R, & Kosslyn, S. M. (Eds). *Memory and mind: A festschrift for Gordon H. Bower*, (pp. 281–305).

Gluck M.A., & Bower G.H. (1988a). From conditioning to category learning: an adaptive network model. *Journal of Experimental Psychology: General, 117*(3), 227–247. https://doi.org/10.1037/0096-3445.117.3.227

Gluck, M. A., & Bower, G. H. (1988b). Evaluating an adaptive network model of human learning. *Journal of Memory and Language*, 27(2), 166–195. https://doi.org/10.1016/0749-596X(88)90072-1

Gluck, M. A., Shohamy, D., & Myers, C. (2002). How do people solve the "weather prediction" task?: Individual variability in strategies for probabilistic category learning. *Learning and Memory*, *9*(6), 408–418. https://doi.org/10.1101/lm.45202

Kamin, L. J. (1969). *Predictability, surprise, attention and conditioning*. In B. A. Campbell & R. M. Church (Eds.), Punishment and Aversive Behavior. New York: Appleton-Century-Crofts.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996a). A neostriatal habit learning system in humans. *Science*, *273*(5280), 1399–1402. https://doi.org/10.1126/science.273.5280.1399

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning Memory*, *1*(2), 106–120. https://doi.org/10.1101/lm.1.2.106

Knowlton, B. J., Swerdlow, N. R., Swenson, M., Squire, L. R., Paulsen, J. S., & Butters, N. (1996b). Dissociations within nondeclarative memory in Huntington's disease. *Neuropsychology*, *10*(4), 538–548. https://doi.org/10.1037/0894-4105.10.4.538

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99*, 22–44.

Kruschke, J. K., & Johansen, M. K. (1999). A Model of Probabilistic Category Learning. *Journal of Experimental Psychology: Learning Memory and Cognition*, *25*(5), 1083–1119. https://doi.org/10.1037/0278-7393.25.5.1083

Lagnado, D. A., Newell, B. R., Kahan, S., & Shanks, D. R. (2006). Insight and strategy in multiple-cue learning. *Journal of Experimental Psychology: General*, *135*(2), 162–183. https://doi.org/10.1037/0096-3445.135.2.162

Lindell, M. K. (1976). Cognitive and outcome feedback in multiple-cue probability learning tasks. *Journal of Experimental Psychology: Human Learning and Memory*, *2*(6), 739–745. https://doi.org/10.1037/0278-7393.2.6.739

Little, D. R., & Lewandowsky, S. (2009a). Better Learning With More Error: Probabilistic Feedback Increases Sensitivity to Correlated Cues in Categorization. *Journal of Experimental Psychology: Learning Memory and Cognition*, *35*(4), 1041–1061. https://doi.org/10.1037/a0015902

Little, D. R., & Lewandowsky, S. (2009b). Beyond Nonutilization: Irrelevant Cues Can Gate Learning in Probabilistic Categorization. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(2), 530–550. https://doi.org/10.1037/0096-1523.35.2.530

López, F. J., & Shanks, D. R. (2008). Models of animal learning and their relations to human learning. In R. Sun (Ed.), *The Cambridge handbook of computational psychology* (pp. 589–611). Cambridge University Press. https://doi.org/10.1017/CBO9780511816772.026

Maddox, W. T., Bohil, C. J., & Ing, A. D. (2004). Evidence for a procedural-learning-based system in perceptual category learning. *Psychonomic Bulletin and Review*, *11*(5), 945–952. https://doi.org/10.3758/BF03196726

Marchant, N., Canessa, E., & Chaigneau, S. E. (2022). An Adaptive Linear Filter model of procedural category learning. *Cognitive Processing*, 23(3), 393–405. https://doi.org/10.1007/s10339-022-01094-1

Marchant, N., & Chaigneau, S. E. (2022). On the importance of feedback for categorization: Revisiting category learning experiments using an adaptive filter model. *Journal of Experimental Psychology: Animal Learning and Cognition, 48*(4), 295–306. https://doi.org/10.1037/xan0000339

Meeter, M., Myers, C. E., Shohamy, D., Hopkins, R. O., & Gluck, M. A. (2006). Strategies in probabilistic categorization: Results from a new way of analyzing performance. *Learning and Memory*, *13*(2), 230–239. https://doi.org/10.1101/lm.43006

Miller, G. A. (2003). The cognitive revolution: A historical perspective. In *Trends in Cognitive Sciences*, 7(3), 141–144. https://doi.org/10.1016/S1364-6613(03)00029-9

Newell, B. R., Dunn, J. C., & Kalish, M. (2011). Systems of Category Learning. Fact or Fantasy? In *Psychology of Learning and Motivation - Advances in Research and Theory* (Vol. 54). https://doi.org/10.1016/B978-0-12-385527-5.00006-1

Newell, B. R., Lagnado, D. A., & Shanks, D. R. (2007). Challenging the role of implicit processes in probabilistic category learning. *Psychonomic Bulletin and Review*, *14*(3), 505–511. https://doi.org/10.3758/BF03194098

Nosofsky, R. M., Gluck, M. A., Palmeri, T. J., Mckinley, S. C., & Glauthier, P. (1994). Comparing modes of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). Memory & Cognition, 22(3), 352–369. https://doi.org/10.3758/BF03200862

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press: NY.

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of comparative and physiological psychology*, *66*(1), 1–5. https://doi.org/10.1037/h0025984

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: current research and theory*. New york: Appleton-Century-Crofts.

Rolison, J. J., Evans, J. S. B. T., Walsh, C. R., & Dennis, I. (2011). The role of working memory capacity in multiple-cue probability learning. *Quarterly Journal of Experimental Psychology*, *64*(8), 1494–1514. https://doi.org/10.1080/17470218.2011.559586

Schultz, W. (1999). The reward signal of midbrain dopamine neurons. *Physiology*, 14(6), 249–254. https://doi.org/10.1152/physiologyonline.1999.14.6.249

Sewell, D. K., Warren, H. A., Rosenblatt, D., Bennett, D., Lyons, M., & Bode, S. (2018). Feedback Discounting in Probabilistic Categorization: Converging Evidence from EEG and Cognitive Modeling. *Computational Brain & Behavior*, *1*(2), 165–183. https://doi.org/10.1007/s42113-018-0012-6

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A Re-Examination of Probability Matching and Rational Choice. *Journal of Behavioral Decision Making*, *15*(3), 233–250. https://doi.org/10.1002/bdm.413

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*(1), 3–22. https://doi.org/10.1037/0033-2909.119.1.3

Smith, E. E., & Grossman, M. (2008). Multiple systems of category learning. *Neuroscience and Biobehavioral Reviews*, *32*(2), 249–264. https://doi.org/10.1016/j.neubiorev.2007.07.009

Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 88(2), 135–170. https://doi.org/10.1037/0033-295X.88.2.135

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*(7), 309–318. https://doi.org/10.1016/j.tics.2006.05.009

Thomas, M. S. C., & McClelland, J. L. (2008). Connectionist models of cognition. In R. Sun (Ed.), *The Cambridge handbook of computational psychology* (pp. 23–58). Cambridge University Press. https://doi.org/10.1017/CBO9780511816772.005

Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, *185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Wilson, R. C., & Collins, A. G. E. (2019). Ten simple rules for the computational modeling of behavioral data. *ELife*, *8*, 1–33. https://doi.org/10.7554/eLife.49547