

An Associative Model of Geometry Learning: A Modified Choice Rule

Noam Y. Miller and Sara J. Shettleworth
University of Toronto

In a recent article, the authors (Miller & Shettleworth, 2007) showed how the apparently exceptional features of behavior in geometry learning (“reorientation”) experiments can be modeled by assuming that geometric and other features at given locations in an arena are learned competitively as in the Rescorla-Wagner model and that the probability of visiting a location is proportional to the total associative strength of cues at that location relative to that of all relevant locations. Reinforced or unreinforced visits to locations drive changes in associative strengths. Dawson, Kelly, Spetch, and Dupuis (2008) have correctly pointed out that at parameter values outside the ranges the authors used to simulate a body of real experiments, our equation for choice probabilities can give impossible and/or wildly fluctuating results. Here, the authors show that a simple modification of the choice rule eliminates this problem while retaining the transparent way in which the model relates spatial choice to competitive associative learning of cue values.

Keywords: spatial learning, geometric module, Rescorla Wagner model, associative learning, conditioned inhibition

When animals learn the location of reward in arenas of various shapes, cues to the shape or geometry of the arena seem to have a special status in that they may not be blocked or overshadowed by other cues such as the colors of walls, even when those cues are better predictors of the reward’s location. Learning about such landmarks or beacons may even be potentiated by geometry (feature enhancement; Miller & Shettleworth, 2007; review in Cheng & Newcombe, 2005). We (Miller & Shettleworth, 2007) recently proposed a simple mathematical model that accounts for many of the puzzling features of such geometry learning (or “reorientation”; Dawson, Kelly, Spetch, & Dupuis, 2008) experiments.

Our model assumes that the learning underlying spatial choice is described by the classic Rescorla-Wagner equation in which all cues (elements, E) at a location (L) compete for associative strength.

$$\Delta V_E = \alpha P_L (1 - V_L). \quad (1)$$

The probability of choosing a location (P_L) reflects the total associative strength of cues at that location relative to the total associative strengths of all locations,

$$P_L = V_L / \Sigma V_L. \quad (2)$$

The apparently special features of geometry learning arise because the animal’s choices determine the contingencies between cues and reward in a dynamic way and because some cues, such as certain shapes of corners, may be shared among locations. Previously learned or salient nongeometric features enhance learning of

geometry because when the animal visits rewarded locations containing those features it learns about co-occurring geometric cues. The model reproduces the results of a substantial number of recent experiments in both dry arenas and water tanks (Miller & Shettleworth, 2007).

The Flaw Identified by Dawson et al.

Dawson et al. (2008) have identified an important mathematical error in our formulation of the model. By making corner choice (P) directly proportionate to associative strength (V_L) and allowing V_L to be negative if the net associative strength of elements at a particular location is inhibitory, the model gives negative values of P_L for certain values of α or after many iterations (i.e., trials). Also, because the sum of the probabilities is always 1, under these circumstances values of P_L for other locations become larger than 1. In addition, when this occurs, values of the various elements reach unreasonable values and fluctuate wildly. As Dawson et al. (2008) point out, this problem can be quite easily remedied by changing the equation for P in such a way that P is guaranteed to remain between 0 and 1. A variety of choice functions may be used including exponential functions (Couvillon & Bitterman, 1985) or as Dawson et al. (2008) suggest the logistic equation or the output activity of a perceptron. Below, we suggest another solution, which also eliminates the problem.

It is important to note that only our choice rule (or performance rule) needs to be modified (Equation 2). The equation that regulates learning in the model (Equation 1) remains unchanged, as do the basic qualities of the model. Whatever performance rule is used to guide choice, so long as choice remains proportional in some way to what has been learned about the elements present at each location, the structure of the model remains the same. A simple modification of the model that satisfies this condition and solves the problem is to set the overall associative strength of a given location (V_L) to 0 if V_L is negative. This is comparable to stating that a location cannot be chosen less often than never, even if it is

Noam Y. Miller and Sara J. Shettleworth, Department of Psychology, University of Toronto.

Correspondence concerning this article should be addressed to Noam Y. Miller, Department of Psychology, University of Toronto, 100 St George Street, Room 4020, Toronto, Ontario, M5S 3G3 Canada. E-mail: noam.miller@utoronto.ca

inhibitory overall. Choices are allocated among the remaining locations according to what can now be thought of as their relative attractiveness, r . Thus we set the attractiveness of a location (r_L) to be the sum of the associative strengths of all the elements at that location (as in our original definition of V_L) if that sum is positive, and set $r_L = 0$ otherwise (i.e., the attractiveness of an inhibitory location is 0). Equation 2 then becomes:

$$P_L = r_L / \sum(r_L). \quad (3)$$

Replacing V_L by r_L ensures that the sum of the choice probabilities over all locations will always be 1 and eliminates the problem uncovered by Dawson et al. (2008) while retaining all the predictive power of the original model. If anything, this modification also increases the model's intuitive psychological realism in that it does not postulate negative choices. Moreover, there is some precedent for giving inhibition special treatment in associative models. Wagner and Rescorla (1972; see also Rescorla, 1979) recognized that making inhibition the precise mathematical opposite of excitation in their model generated some unrealistic predictions, and inhibition has remained problematic for that model (Miller, Barnet, & Graham, 1995).

To firmly establish that the aberrant behavior identified by Dawson et al. (2008) is eliminated by the current modification to

the model, we recreated the example they presented. We reran our simulation of Wall, Botly, Black, and Shettleworth's Experiment 3 (2004; for details of the simulation see Miller & Shettleworth, 2007) as Dawson et al. (2008) did, with all values of alpha set to 0.6, using both the original and the modified model. Figure 1 presents the choice probabilities (P) for both versions of the simulation. The top panel, simulated with the original model, reproduces the impossible results observed by Dawson et al. (2008; see their Figure 1B). The bottom panel shows the same simulation run with the modified model. It may be seen that the new version of the model does not give choice probabilities larger than 1 or smaller than 0. This is also true if the simulation is run for many hundreds of trials (data not shown). In fact, we have not found any condition under which the new model will give impossible results.

We have recalculated the results of all the experiments simulated in our original presentation of the model, both the single choice and the multiple-choice versions. The relative strengths of different elements and the relative percentages of choices of different corners are not changed in any of the simulations, both those presented in detail in the article and those whose results are summarized. In addition, the basic phenomenon of feature enhancement continues to appear as before. In most of the simula-




Figure 1. Choice probabilities for the simulation of Wall et al.'s (2004) Experiment 3. Top panel: run with the original model as presented in Miller and Shettleworth (2007); Bottom panel: run with the modified choice rule presented here. Simulations were run with all alpha values set to 0.6 (see Dawson et al., 2008). Correct = rewarded corner of the rectangular enclosure; Rotational = corner diagonally opposite the rewarded corner; Near/Far = remaining corners.

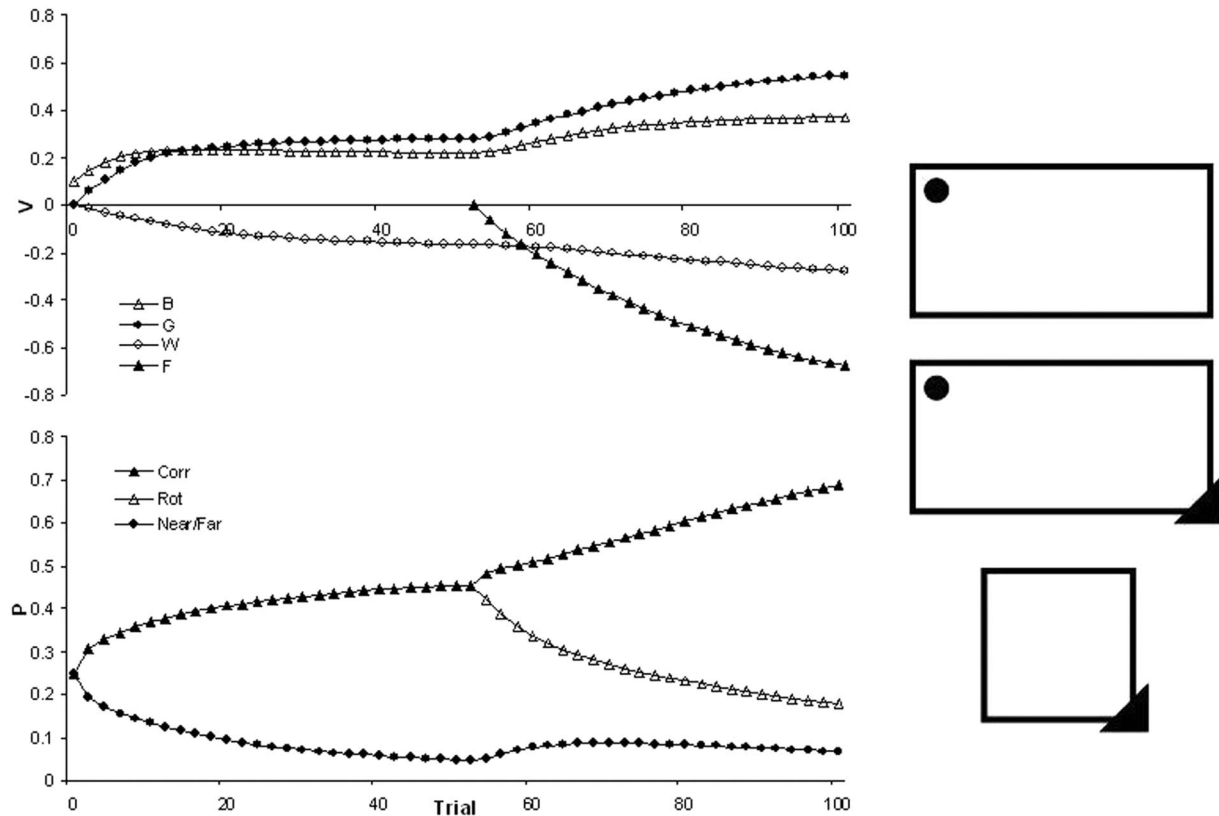


Figure 2. Simulation results of the thought experiment (see text for details). Right panels: diagram of the experiment. The three panels show the three phases of the experiment. The black circle indicates a rewarded location; the black triangle indicates a feature. Left panels: associative strengths (top) and choice probabilities (bottom) for the first two phases of the experiment. Each phase was run for 50 trials (because it took the Model 50 trials of phase 1 to reach the 90% geometrically correct criterion). Upper panel: B = element B as defined in the text; G = geometry of the corners rewarded in Phase I; W = geometry of the corners not rewarded in Phase I; F = the feature. Lower panel: Correct, Rotational, Near/Far as in Figure 1.

tions, the results were identical to those of the original simulation because no location ever became inhibitory overall and the new rule was never invoked.

One simulation in which locations did become inhibitory was that of Cheng's Experiment 3 (1986). Cheng (1986) trained rats to locate a reward in one corner of a rectangular enclosure in which each corner was marked by a distinctive feature. In our simulation of this experiment, the feature at the rotational corner (i.e., the corner diagonally opposite the rewarded corner) acquires a strong inhibitory value because it is paired with the correct geometry, but with no reward. After acquisition, Cheng (1986) tested the rats in a transformed enclosure in which each feature was rotated one corner along. Our simulation assumes that each feature carries its associative strength with it, and that choice of corner during the test is determined by the new total associative strength of each corner. As a result, our original model predicted a negative choice probability for the "far" corner, to which the inhibitory feature from the rotational corner had been moved by the test manipulation (Miller & Shettleworth, 2007, p. 197). Specifically, the original model gave test choice probabilities of 0.32 for both the geometrically correct corners, 0.42 for the near corner, and -0.05 for the far corner. The same simulation run with the modified

model gives 0.3 for the correct and rotational corners, 0.4 for the near corner, and 0 for the far corner. Thus, the modified model does not give negative probabilities while retaining the relative choice percentages among the different corners.

In all of the simulations presented in our original paper, we included an element labeled B, representing contextual cues present at all locations. This element was given an initial associative strength of 0.1 both in order to avoid division by 0 in the first trial's calculation of P and to represent the associative strength resulting from pretraining trials that are common in geometry learning experiments (Miller & Shettleworth, 2007). Since element B is present at all locations, it interacts with all the other elements, both rewarded and unrewarded. In order for a particular location to have a negative associative strength overall, the associative strengths of the inhibitory elements present there would have to be larger, in absolute terms, than sum of the positive terms, which include B. Two related processes usually prevent this from happening:

First, the associative strength of element B itself is decreased by visits to unrewarded corners and increased by visits to rewarded corners. Since most simulations visit rewarded corners far more often than unrewarded corners (which is what the model is sup-

posed to do), V_B tends to either increase overall or to decrease very slowly (as better predictors of reward begin to capture more of the associative strength).

Second, the associative strengths of inhibitory elements decrease only when an unrewarded location is visited. By virtue of being inhibitory, these elements drive the simulation to avoid these locations most of the time, thus retarding the further growth of inhibition. Thus, only very rarely do the conditions exist for inhibitory elements to overcome positive elements that co-occur with them and lead to an aversive location.

To illustrate these dynamic feedback processes in a clear case of inhibitory learning, we model the following thought experiment. Animals are trained initially to choose either of two geometrically identical corners of a rectangular enclosure (Figure 2, top right panel). Reinforcement is given for 50% of correct choices, and the animals are trained until they are about 90% correct (i.e., they choose each of the indistinguishable corners with the correct geometry on 45% of trials). Then, a very salient feature is added to one of the geometrically correct corners and reward is no longer given there, while the unmarked rotational corner with the same geometry always has reward (center right panel). Under these conditions, the feature should become strongly inhibitory while the unmarked opposite corner is eventually chosen most of the time. A test of the inhibitory value of the feature is now conducted by placing the animals in a square enclosure with the feature in one corner (bottom right panel). Here, the animals would be expected to avoid the marked, inhibitory, corner and distribute their choices evenly among the other three. The left panels of Figure 2 show the associative strengths of the various elements (top) and the corner choice probabilities (bottom) for the first two phases of this simulation. The graphs represent the simulations for both the original version of the model and the modified version since, as described above, the modified rule was never invoked during training in this simulation. However, during the test phase, when geometric information is removed, the strongly inhibitory feature causes negative choice probabilities to be predicted by the original version of the model. Specifically, the original model gives choice probabilities at test of -0.41 for the rotational corner (that contains the inhibitory feature), and 0.47 for all other corners, whereas the modified model gives 0 for the rotational corner and 0.33 for the remaining three corners. As with the Cheng (1986) example above, the modified model solves the problem.

In conclusion, the correction to our model proposed here prevents the model from taking on choice probabilities that are greater than 1 or are negative, without otherwise affecting its behavior. While in some extreme cases the rate of learning in the model may vary slightly, there is no difference in the relative associative strengths that it predicts for all the experiments we have simulated

so far. Our model has the advantage of being easily understood on an intuitive level and the dependence of choice on associative strength being obvious. If anything, the modification described here makes the model even more intuitive and follows a tradition in which inhibitory cues seem to require special treatment in models of associative learning. The model shows in a transparent way how what appeared to be exceptional kinds of cue interactions in geometry learning experiments can arise from an unexceptional competition for learning among geometric and other cues. If it turns out that alternative formulations, such as the perceptron proposed by Dawson et al. (2008) and the view-matching process proposed by Cheung, Sturzl, Zeil, and Cheng (2008), can reproduce the same range of findings as our model, a challenge for the future will be to look for ways in which their predictions differ.

References

- Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition*, *23*, 149–178.
- Cheng, K., & Newcombe, N. S. (2005). Is there a geometric module for spatial orientation? Squaring theory and evidence. *Psychonomic Bulletin & Review*, *12*, 1–23.
- Cheung, A., Sturzl, W., Zeil, J., & Cheng, K. (2008). The information content of panoramic images II: View-based navigation in nonrectangular experimental arenas. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*, 15–30.
- Couvillon, P. A., & Bitterman, M. E. (1985). Analysis of choice in honeybees. *Animal Learning and Behavior*, *13*, 246–252.
- Dawson, M. R. W., Kelly, D. M., Spetch, M. L., & Dupuis, B. (2008). Learning about environmental geometry: A flaw in Miller and Shettleworth's (2007) operant model. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*, 415–418.
- Miller, N. Y., & Shettleworth, S. J. (2007). Learning about environmental geometry: An associative model. *Journal of Experimental Psychology: Animal Behavior Processes*, *33*, 191–212.
- Miller, R. R., Barnet, R. C., & Grahame, N. J. (1995). Assessment of the Rescorla-Wagner model. *Psychological Bulletin*, *117*, 363–386.
- Rescorla, R. A. (1979). Conditioned inhibition and extinction. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation: A memorial volume to Jerzy Konorski* (pp. 83–110). Hillsdale: Lawrence Erlbaum Associates.
- Wagner, A. R., & Rescorla, R. A. (1972). Inhibition in Pavlovian conditioning: Application of a theory. In R. A. Boakes & M. S. Halliday (Eds.) *Inhibition and learning* (pp. 301–336). London: Academic Press.
- Wall, P. L., Botly, L. C. P., Black, C. K., & Shettleworth, S. J. (2004). The geometric module in the rat: Independence of shape and feature learning in a food finding task. *Learning & Behavior*, *32*, 289–298.

Received October 30, 2007

Revision received January 4, 2008

Accepted January 4, 2008 ■