

# Get out of the corner: Inhibition and the effect of location type and number on perceptron and human reorientation

Brian Dupuis · Michael R. W. Dawson

Published online: 25 May 2013  
© Psychonomic Society, Inc. 2013

**Abstract** Spatial learning and navigation have frequently been investigated using a *reorientation task* paradigm (Cheng, Cognition, 23(2), 149-78, 1986). However, implementing this task typically involves making tacit assumptions about the nature of spatial information. This has important theoretical consequences: Theories of reorientation typically focus on angles at corners as *geometric cues* and ignore information present at noncorner locations. We present a neural network model of reorientation that challenges these assumptions and use this model to generate predictions in a novel variant of the reorientation task. We test these predictions against human behavior in a virtual environment. Networks and humans alike exhibit reorientation behavior even when goal locations are not present at corners. Our simulated and our experimental results suggest that angles are processed in a manner more similar to features, acting as a *focal point* for reorientation, and that the mechanisms governing reorientation behavior may be inhibitory rather than excitatory.

**Keywords** Spatial learning · Associative learning · Reorientation task · Neural networks · Synthetic psychology

## Introduction


Any mobile agent, capable of autonomously navigating through its world, must be able to find its bearing and orientation to do so. Researchers have developed a number of paradigms to investigate this ability, the foremost of which is the *reorientation task* (Cheng, 1986). In a reorientation task, agents are placed inside a

controlled arena that contains a specific set of cues and are trained to search a particular location for reinforcement. Following testing, the arena is reconfigured, and the changes in agent search patterns are recorded.

A typical reorientation task uses a quadrilateral arena, usually rectangular, with four locations of interest (one in each corner). While some exceptions exist (Cheng, 1986; Newcombe, Ratliff, Shallcross, & Twyman, 2010), the overwhelming majority of reorientation experiments conform to this structure (Cheng & Newcombe, 2005). The emphasis on corners has some desirable properties; for instance, in a rectangle, corners that are diagonally opposite are geometrically identical. Observe, in Fig. 1a, that Locations 1 and 3 have an identical geometry (long wall on the left, short wall on the right, both walls joined at a 90° angle). Because of this, if an agent were choosing a location on the basis of geometry alone, it would have the same likelihood of visiting Location 1 as it would of visiting Location 3. This led to the discovery that rats processed geometric cues even when they could be completely ignored (Cheng, 1986; Gallistel, 1990). In a typical study (for instance, Wall, Botly, Black, & Shettleworth, 2004), Location 1 in Fig. 1a would be reinforced when visited, to indicate that it was the “correct” location. Furthermore, a unique landmark—a nongeometric identifier—would be placed at that location. This landmark provides sufficient information for an animal to learn the “correct” location; geometric information is not required and is, in fact, less reliable, because the geometric information at Location 1 is also present at the nonreinforced Location 3. Nevertheless, if, after training, the rat is placed in a new arena in which the “correct” landmark has been moved to Location 2, its behavior will typically indicate that geometric cues were encoded. That is, it will have a high likelihood of visiting Location 2 in the new arena (indicating that nongeometric cues were learned) but will also visit Locations 1 and 3 (indicating that geometric cues were also learned). Such results led to the development of the geometric module theory

---

B. Dupuis (✉) · M. R. W. Dawson  
Department of Psychology, University of Alberta, Edmonton, AB  
T6G 2E9, Canada  
e-mail: bdupuis@ualberta.ca



**Fig. 1** Possible configurations of targets in a rectangular arena. Many tasks concern themselves solely with corners (a), but similar relationships (i.e., long wall to the left of the target, short wall to the right of the target) can occur in other locations, such as along walls (b). Such arenas therefore need not be limited to four locations of interest (c)

(Cheng, 1986; Gallistel, 1990), which is based on the assumption that geometric cues are processed independently of nongeometric cues and that the processing of geometric cues is mandatory.

The geometric module theory is an example of an insight into navigation provided by the reorientation task. This fairly straightforward task has been a fruitful source of information about navigation and has been used to study a wide variety of organisms, including ants, fish, rats, birds, and humans (review in Cheng & Newcombe, 2005; Cheng, 2008). While this informative task is straightforward to describe and provides data that are easily analyzed, it is important to realize that it is expensive to conduct. Particularly when using animal subjects, an experiment requires considerable commitment of resources, because subjects must be run individually and it takes a fair amount of training for a subject to learn the “correct” location before being placed in a novel arena. It would be convenient if there was a less expensive medium in which to explore the reorientation task, with the aim of discovering interesting hypotheses that could then be tested with a traditional (and more expensive) experiment.

Computer simulations are one less-expensive medium that can be used to explore domains of interest. Lewandowsky (1993) has pointed out that computer simulations can provide several advantages for theory development in cognitive science. These include formalizing a theory in such a way that rigor is improved and providing more precise tools for studying concepts of interest. Additionally, implementing a working computer simulation can be used to reveal tacit assumptions hidden within a theory. Finally, a computer simulation can itself lead to serendipitous findings, particularly when it is presented with novel situations. It would seem that if one had a plausible computer simulation for the reorientation task, it could be used to explore new situations with ease and could possibly generate unexpected predictions. These new simulation-based predictions could then be tested with a traditional

experiment, particularly if a researcher felt that the predictions were interesting enough to warrant the experiment’s expense.<sup>1</sup>

Fortunately, a plausible computer simulation for reorientation has been proposed (Dawson, Kelly, Spetch, & Dupuis, 2010) in the form of a simple artificial neural network called a *perceptron* (Rosenblatt, 1962). In its standard form, a perceptron consists of a single bank of input units that numerically encode patterns of stimuli; these input units are linked via weighted connections to an output unit, which transforms this weighted net input signal into response behavior. The strength of the connection weights is then updated following a specified learning rule, designed to minimize the difference between the output unit’s activity and the desired response to that particular input pattern. A perceptron trained with a standard learning algorithm has been shown to generate most of the interesting regularities found in reorientation task behavior (Dawson, Kelly, et al., 2010). An *operant perceptron* model of reorientation, which uses a more psychologically plausible learning algorithm (in which the perceptron has a chance of not visiting a location, based on the total associative strength at that location, and connection weights are adjusted only when the perceptron chooses to investigate a location), has also been proposed and has shown some promising results (Dawson, Dupuis, Spetch, & Kelly, 2009; Dupuis & Dawson, *in press*). The purpose of the present article is to illustrate how an operant perceptron can be used to explore reorientation by observing the model’s behavior when novel reorientation paradigms are simulated. We demonstrate that this kind of computer simulation can generate interesting predictions that can then be tested using more traditional experimental methodologies.

Part of the power of computer simulations is that, by revealing tacit assumptions about the processes responsible for phenomena of interest (Lewandowsky, 1993), they also provide the means to challenge, or even negate, those assumptions. In the case of reorientation, while the geometric module theory had a strong early impact for many years, it has recently been questioned, with some researchers arguing for it to be abandoned completely (Cheng, 2008; Twyman & Newcombe, 2010). If reorientation is accomplished without the use of a geometric module, what mechanisms might instead be responsible? One alternative to the geometric module is an appeal to the general principles of associationist learning (Miller & Shettleworth, 2007). According to this view, there is no geometric module, but both geometric and nongeometric information are treated in the same manner as being valid cues. Agents use standard learning procedures to associate the various available cues (both geometric and nongeometric) at a

<sup>1</sup> Many of these merits are shared with robotic models of animal behavior. We refer the interested reader to Webb (2000) and Dawson, Dupuis, and Wilson (2010) for more information.

location with the likelihood of being rewarded at that location. A model employing this approach has been shown to be capable of simulating many reorientation task regularities without an appeal to the geometric module (Miller, 2009; Miller & Shettleworth, 2007, 2008). However, serious empirical and theoretical problems with this model have been identified (Dawson, Kelly, Spetch, & Dupuis, 2008; Dupuis & Dawson, *in press*). The standard perceptron avoids these problems and also models reorientation regularities (Dawson, Kelly, et al., 2010). The operant perceptron captures the same regularities as the standard perceptron but does so with a more realistic conception of learning in the reorientation paradigm (Dupuis & Dawson, *in press*). Despite their mathematical differences, both the purely associationist and perceptron models were developed in the spirit of challenging the assumptions about the processes underlying reorientation and as reactions against geometric modularity.

#### Exploring assumptions used to define the reorientation task

Because perceptrons are plausible computer simulations for studying reorientation, they enable the exploration of other challenges, concerning not just assumptions about reorientation processes, but also assumptions about the reorientation paradigm itself. For example, although the prototypical reorientation paradigm (Cheng, 1986) employed a grid spanning the entire arena for measurement, it was noted earlier that a common feature of most modern reorientation experiments is the emphasis on corners as the only locations of interest. Might performance on the reorientation task be affected if agents were trained to go to locations that are not at corners? Are corner locations special in some way?

According to associationist theories of reorientation, corner locations should not be special. These theories, based on the work of Rescorla and Wagner (1972), posit a view of learning largely based on cue competition. Applied in to the reorientation task, this suggests that a location is merely a collection of cues that can be exploited as signals of potential reinforcement. From this perspective, there should be no fundamental behavioral difference between a location of interest at a 90° corner and one along a 180° wall. In effect, a location under consideration along a single wall “divides” the wall into a left and a right segment, exactly as a corner does, with only the angle of intersection distinguishing them. Under this associative viewpoint, it would appear that locations of interest need not be constrained to corners.

Another tacit assumption that guides experimental studies of reorientation concerns the number of target locations. The vast majority of reorientation studies have used quadrilateral arenas—typically, rectangles or squares (review in Cheng & Newcombe, 2005; Cheng, 2008) and, less commonly, in kites (Dawson, Kelly, et al., 2010; Pearce, Good, Jones, &

McGregor, 2004), or parallelograms (Lubyk & Spetch, 2012; Tommasi & Polli, 2004). In these studies, the corners of these arenas have been used as target locations (even if noncorner information is part of the study; e.g., Ratliff & Newcombe, 2008), and therefore, these experiments have studied reorientation using four different locations. Furthermore, nonkite quadrilateral arenas make available only two different instances of geometric cues (long wall on the left and short wall on the right of a corner, short wall on the left and long wall on the right of a corner). Only a handful of studies have used arenas that are not quadrilaterals (i.e., Newcombe et al., 2010; Sturz & Bodily, 2011) and have, therefore, made available more than four potential target locations (again assuming that locations of interest are always positioned at arena corners). To our knowledge, no experimental studies of reorientation have explicitly compared situations in which the number of target locations has been systematically varied.

However, it is important to study the effect of varying the number of target locations, because some theories predict that this variable should affect learning in the reorientation task. For example, Miller and Shettleworth’s (2007, 2008) associative model scales the rate of learning by a measure of the probability of an agent visiting a particular location in its learning equations. This measure is expressed as the net attractiveness of the location as a proportion of the total net attractiveness of every location. This proportion will obviously be affected by the number of locations that are summed in its denominator. All else being equal, this will predict (at least initially) one half the normal rate of learning for tasks with eight locations, relative to tasks with four locations.<sup>2</sup> That is, the Miller–Shettleworth model predicts that learning the reorientation task will slow down as the number of possible locations of interest increases. Do other models, such as the perceptron, also make this prediction?

Furthermore, as more locations are added to a standard reorientation arena (Fig. 1c), a greater variety of geometric cues must be processed. For instance, there are four different geometric configurations that can be distinguished in Fig. 1c (as opposed to two in Fig. 1a and b), which again would be expected to slow reorientation learning in a theory like Miller and Shettleworth’s (2007, 2008) associative model. Does varying the number of target locations, and the number of possible geometric configurations, affect human learning in the reorientation paradigm, and does it do so in a fashion predicted by computer simulations? The present article represents an attempt to begin the exploration of such questions.

The purpose of the present work is to investigate the two main issues raised above. First, we attempt to evaluate the role that the nature of a location—at a corner or at a

<sup>2</sup> Due to other mathematical flaws with the Miller and Shettleworth (2007, 2008) model, detailed at length in Dupuis and Dawson (*in press*), we will not dwell further on this model here.

wall—has on reorientation behavior. Second, we investigate the impact that changing the number of locations of interest has on reorientation behavior. These distinctions are depicted in Fig. 1. Figure 1a illustrates the typical position of possible target locations in a standard reorientation task that uses a rectangular arena. Figure 1b provides an analogous arena, but one in which the locations of interest are not found at corners. Figure 1c shows how one can combine the first two arenas into a third that has eight locations of interest, instead of the typical four.

This article proceeds as follows. It begins by using an operant perceptron (Dupuis & Dawson, *in press*) to simulate the reorientation task in the various arenas illustrated in Fig. 1. These simulations are used to make predictions about the effects of type of location and of number of locations on reorientation behavior. The results of these simulations provide two key predictions: (1) Corner locations are not inherently special, and (2) doubling the number of target locations has a negligible effect on the speed at which the model learns to reorient. Underpinning these predictions is evidence that the mechanism at work may be inhibitory rather than excitatory, which has important theoretical implications. Next, we report the results of testing these predictions, using human subjects in a virtual world. Finally, we explore the similarities and differences between the associationist model and the human data. We argue that the operant perceptron is a useful source of predictions that can be supported by experimental data. As a result, the operant perceptron appears to provide a medium in which reorientation can be plausibly explored for the purpose of seeking surprising and interesting results that can later become the focus of traditional experimentation.

## Simulation

From the perspective of theories of reorientation that appeal to a geometric module (Gallistel, 1990), angle information present at a corner is typically viewed as a global, geometric property. However, from an associationist perspective, a “corner” could be perceived as a visually salient “focal point” that serves as a reference, with the angle simply being a (local) feature of that location. The intersection between walls provides a distinct boundary, from which the length of a wall can be measured. For instance, in Fig. 1b, Location 1 sits at the junction of a short wall on its left and a long wall on its right, with an intersection angle of  $180^\circ$ .

With this in mind, we devised a method of representing any location along the edge of an arena that treats angle information as a feature. This representation permitted us to present locations to perceptrons even when these locations were not at a corner in a reorientation arena.

## The perceptron

### *Perceptron reorientation*

As was noted earlier, a perceptron (Rosenblatt, 1962) is a simple connectionist network in which a set of input units are directly connected to an output unit via weighted connections. The input units represent stimuli; their activation causes signals to be sent through the weighted connections to produce a response in the perceptron’s output unit. Feedback can be provided to the network about its response so that it can modify its connection weights. This permits the perceptron to learn to generate a desired response to each stimulus in a set of training patterns.

To simulate the reorientation task, each location of interest in an arena is represented as a stimulus in the set of training patterns. For each of these patterns, input unit activity is used to represent which cues (geometric and nongeometric) are present at a particular location. If a location is deemed to be “correct,” the perceptron is reinforced when that location’s cues are presented. If a location is not deemed to be “correct,” the perceptron is not reinforced when that location’s cues are presented. In other words, the perceptron is trained to produce an activity of 1 to sets of cues corresponding to “correct” locations and an activity of 0 to sets of cues corresponding to “incorrect” locations.

In order to train the perceptron to learn to reorient in a particular arena, one must make design decisions about how to represent the available cues and about the learning rule that is used to modify the network’s connection weights. The details of these design decisions are provided below.

### *Defining the task: Stimuli*

Each location identified in Fig. 1 can be defined as a collection of properties, which are presented to the perceptron as a pattern of unary-coded inputs. That is, each of the perceptron’s input units encodes the presence or absence of a specific cue. Each of these units is turned on (activated with a value of 1) when the property it encodes is present and is turned off (activated with a value of 0) when that property is absent. In the present simulation, each location of interest is defined by three types of cues: the length of the walls on either side of the location, the angle between the walls where they join, and the kind of local landmark that can be present at the location. Seventeen different input units were used to represent the possible values of these cues, as is summarized in Table 1.

The angle units (1–2) identify the angle of intersection at the location. These units have one value for locations at corners ( $90^\circ$  angle) and another value for locations along walls ( $180^\circ$  angle). The feature units (3–11) represent the collection of nongeometric properties present at a given location. For parsimony with the experiments described

later in the article, these units are named after colors; as such, these units can be thought of as representing the color of an object at the location.

The length configuration units (12–17) represent the specific set of wall length properties present at a location. For example, one unit is turned on for a location at the intersection of a wall of length three with a wall of length six, while another might be turned on if the location lies between walls of length two and one. This is an extension of Miller and Shettleworth's (2007, 2008) representation for specific geometries that allows for a number of possible configurations—up to six in the current simulation. This is required when more than four locations of interest are used (Fig. 1c).

When put together, this encoding can represent any possible location of interest as a string of 0 s and 1 s reflecting the absence or presence of the corresponding cue at that location. For instance, if Location 1 in Fig. 1b (180°, two-length wall to the left, four-length wall to the right) contained a “blue” feature, the location would be presented to the perceptron as “1000010000000100.”

This particular set of design decisions defines this encoding as a purely local code: Each pattern contains only information present at the location it represents and no information from any other location. Similarly, this encoding contains no global representation of the arena, either explicitly (i.e., a principal axis, Cheng & Gallistel, 2005; or relative dimensions, Huttenlocher, 2003) or implicitly (as in Miller & Shettleworth's [2007, 2008] model summing across all locations), save for the number of patterns presented. Indeed, the perceptron is unable to distinguish to which arena a duplicated pattern belongs; the above example location code would also be seen when Location 1 in Fig. 1c was presented, for instance, since this location is geometrically and featurally identical to Location 1 in Fig. 1b. The only information available to the perceptron at any given time is found in the cues present at the location under consideration; in order to detect additional unrepresented features in this encoding, the perceptron would require hidden units (Dawson, 2004; Rummelhart, Hinton, & Williams, 1986) or a similar architectural adjustment.

#### *Defining the task: Response*

This simulation includes a single output unit that uses the logistic activation function (Dawson, 2008) to convert the total weighted signal coming from the input units into a response that can range between 0 and 1. For locations that are reinforced, the perceptron is trained to turn on (output activity=1); for locations that are not reinforced, the perceptron is trained to turn off (output activity=0). Because, during learning, perceptron activity falls in the

continuous range between 0 and 1, at any moment in time, the perceptron's output can be interpreted as its estimation of the conditional probability of reinforcement at a location, given the cues at that location (Dawson et al., 2009; Dupuis & Dawson, *in press*).

#### *Training method*

A perceptron's response to particular patterns of stimuli is not perfect; each generated response differs from a desired response by some error amount. This error is then used by a learning rule to adjust the perceptron's connection weights such that subsequent presentations of that pattern of stimuli produce a smaller error. Here, we employ the gradient-descent learning rule (Dawson, 2004, 2008), which has desirable properties when working a logistic perceptron response.

This output response provides a critical distinction between neural network models and traditional associative models in the style of Rescorla and Wagner (1972). The perceptron's output activity allows it to convert *associative strength* of assorted cues into a model of *behavior*. This stage is absent from traditional associative models, which describe only the indirectly observable associative strength. Not only does this difference allow perceptrons to produce different predictions from formally equivalent associative models (Dawson, 2008), but also it allows us to adjust the model's learning to reflect different patterns of behavior.

In a standard perceptron model, the connection weights would be updated after presenting any location to the network; that is, there would be no model of choice behavior during learning, much like the definition of classical conditioning. In the present model, we use network output as a measure of behavior to adjust the perceptron's learning from classical conditioning to operant conditioning, where it is allowed to “choose” whether or not to investigate a particular location, and this investigation (rather than rote presentation) governs its learning. Instead of updating connection weights after every pattern of cues is presented, the perceptron's output response to that pattern is used as the *conditional probability* of updating weights on this presentation, given the cues presented. With each presentation, a random number between 0 and 1 is generated and compared with the output response; if the random number exceeds the output response, the connection weights are not updated, and the next pattern is presented. In effect, the perceptron will choose whether or not to visit a location with a probability based on how attractive the cues at that location are, and it will learn only from locations it chooses to visit. This algorithm is detailed at length in Dupuis and Dawson (*in press*) and in a more abbreviated form in Dawson et al. (2009).

**Table 1** Encoding of a location's properties and the agent's response using an operant perceptron

Unit	Codes For	Encoding	Example Values	Example Encoding
Inputs 1–2	Angle at target	Unary coding, 2 units	90°	1 0
			180°	0 1
Inputs 3–11	Color of target	Unary coding, 9 units	Red	1 0 0 0 0 0 0 0
			Green	0 1 0 0 0 0 0 0
			Yellow	0 0 1 0 0 0 0 0
			Blue	0 0 0 1 0 0 0 0
			Orange	0 0 0 0 1 0 0 0
			Black	0 0 0 0 0 1 0 0
			Brown	0 0 0 0 0 0 1 0
			Purple	0 0 0 0 0 0 0 1
			White	0 0 0 0 0 0 0 1
Inputs 12–17	Configuration of wall lengths at target	Unary coding, 6 units	Left 3, Right 6	1 0 0 0 0 0
			Left 6, Right 3	0 1 0 0 0 0
			Left 2, Right 1	0 0 1 0 0 0
			Left 2, Right 4	0 0 0 1 0 0
			Left 4, Right 2	0 0 0 0 1 0
			Left 1, Right 2	0 0 0 0 0 1

### Simulation specification

**Training** The present simulation includes two experimental conditions: one with four locations of interest, and one with eight locations of interest. Each location is present at either a wall or a corner and contains a unique feature cue (i.e., a colored object). Since there can be up to eight possible locations within one condition, this demands eight unique feature cues. Within each condition, networks are trained to investigate just one location; this location is reinforced, while all others are not reinforced. The reinforced location could be present at a wall or a corner, producing a 2 (four-vs.-eight)  $\times$  2 (corner-vs.-wall) design.

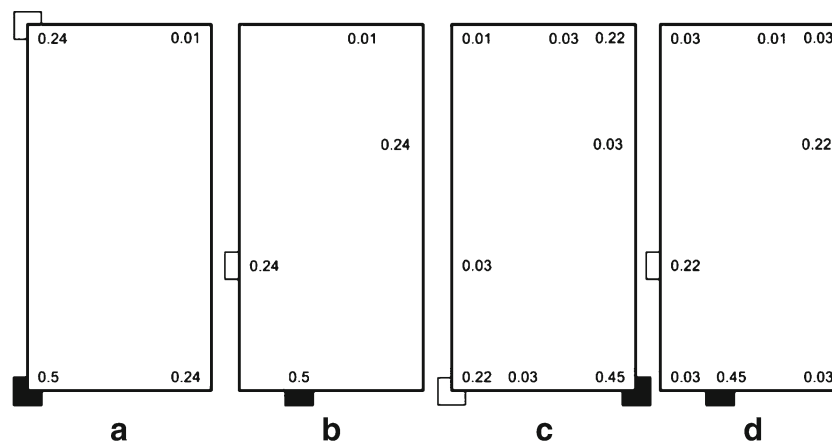
All networks were initialized with all biases and connection weights equal to zero and were trained with a learning rate of 0.1. Five networks in each condition were trained to convergence. For counterbalancing, two possible reinforcement locations were used in each simulation; for example, in the four-location, corner-goal task (Fig. 1a), one group of networks is reinforced at Location 1, while another is reinforced at Location 2. No appreciable difference was found between these groups within a particular task, so their results are reported together here. (That is, each value is averaged from 10 networks.)

**Testing** Testing the perceptron involves presenting patterns of cues corresponding to transformed arenas and measuring the perceptron's output response to these novel patterns. Due to the operant nature of its training algorithm, the perceptron's output response is both its estimation of the conditional probability of reward at the location given the cues at that location and its likelihood of choosing to visit that location.

There are two types of transformed arenas common to reorientation studies: *affine* transformations and “*featureless*” transformations. Affine transformations place feature cues and geometry cues in conflict with each other: a chosen location could be consistent with the geometry present during training or the features present during training. For instance, in Fig. 1a, a subject might find reinforcement at Location 1, along with a unique feature. When placed in an arena with an affine transformation, that unique feature might now be present at Location 2. Location 2 is consistent with training in terms of features, while Locations 1 and 3 are consistent with training in terms of geometry. Meanwhile, a featureless transformation replaces all unique feature cues with indistinguishable ones, forcing the model to base its decisions solely on encoded geometry. In the present simulation, we perform a featureless transform by simply turning off all feature units that were present during training and activating a novel “white” feature unit in their place.

With four locations, we can also test for generalization across angle cues by observing a corner-trained network's response to a wall-locations-only arena (that is, a network trained in Fig. 1a but tested in Fig. 1b) and vice versa. In this scenario, each location now appears with novel angle and length configuration cues, as opposed to an affine transformation, which has novel length configurations but consistent angles. Since these two conditions do not share exact wall lengths, no choice can be consistent with wall length geometry from training.

With eighth locations, one can also do a *partial* transformation. While affine-transformed arenas have consistent angle information (targets that were present at corners are present at



**Fig. 2** Example perceptron response frequencies in affine-transformed arenas. The white boxes indicate where features (and reinforcement) were found during training; the black boxes indicate their locations after an affine transformation. Perceptrons were reinforced in corners (**a** and **c**) or walls (**b** and **d**), and with four (**a** and **b**) or eight (**c** and **d**) total locations

corners during testing), a partial transform places them in conflict. Both transformations have novel wall length geometries, as compared with the training condition.

Following training, each network was presented with probe trials in three transformed arenas. For four locations, these were affine, generalized, and featureless arenas. For eight locations, these were affine, partial, and featureless arenas. Each network's output responses were recorded for each of these locations; these responses were averaged across the five networks present at each condition.

In the reorientation task literature, it is common to report responses in terms of the frequency with which each location is chosen. However, the perceptron responds to each location individually, producing the probability of choosing to act at that specific location; these probabilities need not sum to 1 across all locations within an arena. In order to convert the former into the latter, we divided the response to a specific location by the sum of responses to all locations within a given arena; this method has previously been used to successfully predict several key reorientation behavior regularities (Dawson, Kelly, et al., 2010).

## Results

Across all conditions, networks converged after an average of 4,810 presentations of the training set (a single presentation of each pattern [location] in the training set in a random order is called a *sweep*), with the fastest training occurring after 4,614 sweeps of training and the slowest training requiring 4,973 sweeps. Due to the perceptron's specified learning rate parameter of 0.1, training times of this magnitude are not uncommon for problems of this size; the more pertinent observation is that the range of training times is quite narrow and is not significantly different across all training conditions,  $F(3, 36) = 1.95, p = .14$ .

from which to choose. Response frequencies are normalized from perceptron activity in each arena to allow for comparison with extant animal reorientation data, including human experiments reported later in this article. Note the characteristic *rotational error* in all four conditions. Refer to the [Appendix](#) for complete response data in all conditions

## Network responses

The network model's responses to each location in each transformed arena, expressed both as response activity and as choice frequencies, are reported in detail in the tables presented in the [Appendix](#). *Network activity* refers to the activity in the network's output unit given the cues presented at a particular location; this is interpreted as the network's estimate of the conditional probability of reinforcement at a location given the cues present at that location. These conditional probabilities are converted into network *frequency* through normalization within each condition (Dawson, Kelly, et al., 2010). The tables also include a summary of human responses in similar conditions in experiments that were inspired by the simulation results. The human responses in the table are covered in more detail when the human experiments are discussed, below.

The first major prediction generated by the model is that there does not appear to be a significant difference in reorientation behavior between networks trained with locations in corners and networks trained with locations along walls. In conditions with four locations of interest and eight locations of interest, whether reinforced at a corner or at a wall, the perceptron converged after a similar number of sweeps of training. Furthermore, in all cases, the same broad pattern of behavior holds: The perceptron responds most strongly to locations containing the (unique) feature cue present during training, but that cue did not prevent the encoding of either geometric cue. That is, even within the featureless arena, the perceptron still estimates that locations with the same wall length configuration and/or angle amplitude as the training location have a greater likelihood of reward than do locations missing those cues.

Additionally, the perceptron produces characteristic *rotational error* behavior common to reorientation tasks (Cheng,

**Table 2** Operant perceptron connection weights for each simulation

Unit Type		Problem and Type			
		Four		Eight	
		Corner	Wall	Corner	Wall
Angle	Bias	-0.41	-0.40	-1.14	-1.14
	90	-0.41*	0.00	0.00*	-1.14
	180	0.00	-0.40*	-1.14	0.00*
Color	Red	-0.66	-2.30	-0.57	-0.28
	Green	3.19*	-0.67	3.30*	-0.29
	Yellow	-2.29	-0.65	-2.18	-0.28
	Blue	-0.66	3.21*	-0.55	-0.29
	Orange	0.00	0.00	-0.29	-0.56
	Black	0.00	0.00	-0.28	-0.56
	Brown	0.00	0.00	-0.29	-2.18
	Purple	0.00	0.00	-0.28	3.30*
	White	0.00	0.00	0.00	0.00
	Configuration	3 / 6	-1.31	0.00	0.00
6 / 3		0.91*	0.00	0.00	0.00
2 / 1		0.00	-1.32	-0.57	-1.12
2 / 4		0.00	0.91*	-0.57	1.12*
4 / 2		0.00	0.00	1.12*	-0.57
	1 / 2	0.00	0.00	-1.12	-0.57

\*Corresponds to reinforced location

1986); that is, where features and geometry conflict in the same arena, the perceptron responds to the feature more frequently than to any other single location, but taken as a whole, locations with correct geometry are chosen with higher frequency. This pattern appears in both the four-location and eight-location tasks, as illustrated for the affine transformation in Fig. 2 (for more detail on the other conditions, please refer to the Appendix). Furthermore, it occurs even if the angle information changes between conditions; for instance, the generalized arenas in the location task still produce this pattern, even though the exact configuration of geometries present in this condition are novel.

### Connection weights

To understand why these networks behave in this manner, we turn next to their connection weights. Since 10 networks completed each training condition, their connection weights were averaged to produce a summary of how a typical network solved that particular problem. This summary is presented in Table 2.

An examination of this table reveals that, within the four-location task, the bias and reinforced angle units assume negative values, while nonreinforced angle units assume a value of 0. This informs us that, before considering wall length configuration or

feature information, the network initially tends to turn off (output activity and, thus, probability of investigating a location approaching 0) at any given location. In the eight-location task, however, this is slightly different: While the bias remains negative, the reinforced angle assumes a 0 weight, while the nonreinforced angles assume a strong negative weight. Despite this difference, this pattern of weights, in the absence of other cues, produces identical behavior to the location network.

It is only after the network considers other cues that it begins to overcome this negative association and develop a moderate probability of investigating a given location. Within the four-location task, the wall length configuration corresponding to the reinforced location assumes a positive value with magnitude slightly larger than the magnitude of the bias and the angle at that location. A similar result occurs in the eight-location task, where the correct wall length configuration and the bias effectively cancel out and the correct angle has a weight of 0 (the only situation where a cue present during training assumes a 0 weight). In both of these cases, the net input is close to 0; the output unit's logistic function translates this into a .5 probability of acting, given those cues. In other words, for both the four-location and eight-location tasks, if the networks encounter a location with the correct geometry but lacking any feature, they are as likely as not to choose to investigate that location. The overall choice frequency behavior this produces will vary depending on the number of locations (see the tables in the Appendix); however, the underlying mechanism is identical. It is interesting to note that, ignoring features, the "correct" wall length configuration is reinforced on 50 % of its presentations (the reinforced location and its nonreinforced rotational equivalent), while the "incorrect" configurations present in any condition are reinforced 0 % of the time, and the perceptrons' responses converge to match these probabilities. The operant perceptron has already been established to match probabilities in classical choice-behavior tasks (Dawson et al., 2009); for it to exhibit this behavior in a reorientation context reinforces Miller and Shettleworth's (2007) conceptualization of reorientation as an operant task.

The feature cue connection weights tell an unsurprising story in both the four-location and eight-location tasks. The feature that was reinforced during training assumes a very strong positive weight, while the feature rotationally opposite the reinforced location (i.e., the other location with identical geometric cues) assumes an equally strong negative weight. The positive magnitude of the weight given to the correct feature far exceeds the negative value of the bias plus any incorrect geometric cue; that is, the network has a high probability of acting when presented with the correct cue, even if both angle and wall length configuration cues are incorrect. Meanwhile, all other features take on a moderate negative weight. In the context of the geometric cues discussed above, this informs us that the network is inherently hesitant of investigating any location but that the presence of a correct



feature is sufficient to overcome this hesitancy. Furthermore, the feature present at the rotational equivalent of the reinforced location during training assumes a negative strength sufficient to overcome correct geometric cues—in effect, becoming a reliable indicator of *no* reinforcement.

## Discussion

The operant perceptron's behavior on these simulations allows us to generate novel empirical predictions. To begin, the network used the same encoding for all conditions (four locations or eight locations and wall reinforcement or corner reinforcement) and was able to converge in all of these conditions without difficulty with the same amount of training. Therefore, the operant perceptron predicts that similar mechanisms are at work regardless of the global shape of the arena and that changing the number of locations of interest will have a negligible effect on the difficulty of the task. These predictions are broadly compatible with previous empirical work on multiple-location reorientation (Newcombe et al., 2010) but are incompatible with theories that include an implicit representation of the global environment (Miller & Shettleworth, 2007, 2008).

Furthermore, the operant perceptron does not predict any real difference between tasks where the locations of interest are found within corners and tasks where such locations are not found at corners. In both cases, networks were able to learn the task, encoding sufficient geometric cues to reorient and producing comparable behavior when presented with transformed arenas. This behavior persisted even if the cue types were completely novel, suggesting some degree of generalization—although the network predicts that the mechanism behind this generalization is inhibitory.

We can elaborate on this inhibitory mechanism by examining the connection weights in Table 2. Specifically, the networks learned that particular wall length configurations signaled that a location was *not* reinforced and learned that a particular color's rotational opposite was a reliable indicator of no reinforcement. When the networks were presented with the transformed arenas, they did not respond to the novel geometry at all; they had not learned that such configurations signaled no reinforcement. Instead, the network responds at chance values to each location, except for the two locations containing the “correct” feature and its rotational opposite. Rather than developing an explanation of what the agent may be searching for in these cases, a study of connection weights informs us that we might instead be focusing on what the agent is *avoiding*. This tendency to emphasize excitation at the expense of inhibition when explaining learning is a tacit assumption present in many different theories of learning (Rescorla, 1967); the operant perceptron model reinforces this point and reminds us of the need to check such assumptions.

## Experiments

The operant perceptron has generated some interesting predictions on novel permutations of the reorientation task. Specifically, the operant perceptron makes two broad claims: first, that there is no appreciable difference in reorientation behavior among groups trained with locations in corners or along walls, and second, that there is no appreciable change in difficulty when the number of salient locations changes. Do these predictions hold under laboratory conditions with humans? To test these predictions, we conducted a series of basic reorientation experiments using human subjects.

Our experiments are organized into three studies. Study 1 involves two groups of subjects trained on a four-location reorientation task; one group is trained on corner locations, and another trained on wall locations. The locations in these tasks correspond to those shown in Fig. 1a and b.

Study 2 is analogous to Study 1, except that the training arenas have eight locations of interest, as in Fig. 1c.

Immediately after completing Study 1 or Study 2, each participant also completed the task described in the other study; Study 1 participants completed the four-location task and then immediately progressed through the eight-location task exactly as described in Study 2, and vice versa. This allows a direct comparison of the difficulty of reorientation in arenas with four locations and eight locations. Furthermore, this manipulation can test for order effects: Did subjects learn either task faster, and did the first task facilitate learning the second? These comparisons are the focus of Study 3.

### Study 1: Four locations


#### Method

**Subjects** Subjects were 36 University of Alberta undergraduates (30 female), who received course credit for participation. Recruitment criteria required subjects to have normal color vision.

**Apparatus** The environment was constructed using the fAARS-Lite platform (Dupuis, 2012; Gutiérrez, 2012; Lubyk, Dupuis, Gutiérrez, & Spetch, 2012), which simulates first-person 3-D movement in a virtual world. This virtual world contained a number of rectangular arenas (17.2 × 8.6 m), consisting of matte-gray walls and floors with black, visually obvious edges. The walls were high enough to extend beyond the default field of vision in all possible subject locations and orientations.

Subjects arrived in an arena in its center, facing a random direction (one of the eight cardinal or ordinal directions, chosen at random on each arrival). Subjects could move their perspective through these arenas with the arrow keys.

**Fig. 3** Example stimuli from Study 1, from the perspective of a subject standing in the center of a rectangular arena with visually salient locations of interest placed at the corners (a) or the walls (b). Compare with Fig. 1. These locations could have many possible colors and patterns (c), including matte white, although some colors appear only in Study 2



*Stimuli* Attention was called to locations of interest through brightly colored cylinders (1.5-m radius, same height as surrounding walls) placed against the walls of the arena. These locations were placed in two possible configurations, with the locations of interest being set at corners or along walls. These configurations correspond to Fig. 1a and b. During training, these cylinders had one of four colored textures placed over a white background: green checkerboard, red diagonal stripes (upper-left to lower-right), yellow diamonds, or blue horizontal bars, listed in the order in which they appeared when one looked clockwise from the center. Figure 3 presents examples of these stimuli.

During testing, the positioning of these locations shifted into one of three possible configurations: affine, generalizing, or featureless. In the affine condition, each location had been shifted one “slot” clockwise; in Fig. 1a, the green location had been present at Location 1 during training and would be at Location 2 in an affine-transformed arena. In the generalizing condition, the targets were shifted to a novel geometry—that of the other group’s training condition. That is, a subject trained in Fig. 1a would experience Fig. 1b as its generalizing condition. Finally, the featureless transform removed all distinguishing information from the cylinders; in place of their brightly colored patterns, subjects simply saw blank, white pillars.

*General procedures* Participants were pseudorandomly divided into four groups based on two possible categories (counterbalancing for gender). One of these divisions was based on target location: Half the subjects would be trained with locations along walls, and half the subjects would be trained with locations at corners. Similarly, subjects were split into two reinforcement groups before the experiment began: those who would receive reinforcement at a location with a long wall on its left (group A) and those who would receive reinforcement at a location with a long wall on its right (group B). The division between group A and group B also counterbalances for distance to first location (where applicable); for instance, in the wall group (see Fig. 1b), half the subjects would be reinforced at

Location 1, which is close to the start location, and half would be reinforced at Location 4, which is further away.

Upon arrival, subjects were instructed on how to move around in the virtual world and were given time in a “welcome” room (a curved hallway with arrows pointing to a door at its end) to practice movement before the experiment began. Instructions were given to find a “correct location” inside each new room they saw; these instructions deliberately avoided the words “corner” and “wall.” Subjects were told that they would see a display informing them whether their choice was correct or incorrect. Occasionally, they were told, the display would say “no feedback” regardless of the accuracy of their choice. To encourage a consistent strategy over time, subjects were told that they would be awarded points for correct choices (even if the display said “no feedback”) and that they should maximize their score.

*Training* Training consisted of blocks of 10 presentations of the training arena. Each presentation allowed the subject any amount of time to move freely about the enclosure but ended when they moved into one of the target locations, receiving feedback as described above (lasting for 3 s) before appearing in the center of the training enclosure again. Training continued until the participant made eight “correct” choices in a single block (if, at the end of a block, they had failed to make eight “correct” choices, they restarted the block instead), at which point they progressed into a nonreinforced training phase. During this phase, subjects had a 50 % chance of seeing a “no feedback” message after making a choice. After making eight correct choices in a single block during this phase (those scoring less instead restarted the nonreinforced training block), testing began.

*Testing* Subjects received five test trials in each of the three transformation conditions (affine, generalized, and featureless). The test trials were presented in random order. Subjects were aware that their choices would always receive a “no-feedback” response during this phase.

Following testing, subjects completed a posttesting retention test—10 no-feedback trials in their original training enclosure. Subjects must have made seven correct choices during this test to be included in analysis. Following this test, subjects were *not* immediately debriefed; instead, they proceeded to Study 3.

### Results and discussion

Five subjects failed the posttesting retention test, resulting in 31 subjects (26 female) included in analysis. Early analysis indicated that the data do not conform to normal distributions. Therefore, subjects' choice data were analyzed using randomization tests, employing bootstrapping methods to obtain confidence intervals on Cohen's  $d'$  measure of effect size (Edgington, 1995; Efron & Tibshirani, 1994).

This analysis showed no difference in number of training blocks between the genders (males,  $M = 2.0$ ,  $SD = 0.71$ ; females,  $M = 1.88$ ,  $SD = 1.34$ ;  $d' = 0.1$ , 95 % CI  $[-0.56, 1.32]$ ,  $p = .86$ ) or between different reinforcement groups (group A,  $M = 1.6$ ,  $SD = 0.63$ ; group B,  $M = 2.2$ ,  $SD = 1.6$ ;  $d' = -0.48$ , 95 % CI  $[-1.03, 0.24]$ ,  $p = .19$ ). Similarly, subjects did not exhibit any significant differences in the amount of blocks required to learn the task if they were reinforced at walls or at corners (walls,  $M = 2.25$ ,  $SD = 1.61$ ; corners,  $M = 1.53$ ,  $SD = 0.52$ ;  $d' = 0.60$ , 95 % CI  $[-0.02, 1.11]$ ,  $p = .08$ ).

Within the affine test arena, there were no significant differences between groups in terms of the proportion of choices made that were geometrically consistent (wall,  $M = 0.22$ ,  $SD = 0.29$ ; corner,  $M = 0.09$ ,  $SD = 0.14$ ;  $d' = 0.58$ , 95 % CI  $[-0.09, 1.15]$ ,  $p = .11$ ) or featurally consistent (wall,  $M = 0.74$ ,  $SD = 0.33$ ; corner,  $M = 0.89$ ,  $SD = 0.18$ ;  $d' = -0.58$ , 95 % CI  $[-1.18, 0.03]$ ,  $p = .11$ ) with training. There were also no significant differences between groups in terms of featurally consistent choices within the generalized arena (wall,  $M = 0.91$ ,  $SD = 0.16$ ; corner,  $M = 0.87$ ,  $SD = 0.25$ ;  $d' = 0.17$ , 95 % CI  $[-0.53, 0.76]$ ,  $p = .63$ ). Finally, when subjects were tested in featureless arenas, both groups made similar amounts of geometrically consistent choices (wall,  $M = 0.58$ ,  $SD = 0.35$ ; corner,  $M = 0.57$ ,  $SD = 0.34$ ;  $d' = 0.03$ , 95 % CI  $[-0.59, 0.77]$ ,  $p = .96$ ).

Since neither corner-reinforced nor wall-reinforced subjects showed any differences in choice behavior, they were pooled together to test whether their geometric or feature choices were significantly different from chance. Within the affine arena, subjects' choices followed the feature at a rate significantly higher than chance (95 % CI  $[0.73, 0.90]$ ,  $p < .05$ ), and they made choices consistent with training geometry significantly less often than chance (95 % CI  $[0.08, 0.24]$ ,  $p < .05$ ). Feature-consistent choices were made in the generalized arenas more often than chance (95 % CI  $[0.82, 0.95]$ ,  $p < .05$ ). Within the featureless arena, however, subjects' choices did not significantly differ from chance (95 % CI  $[0.46, 0.68]$ ,  $p > .05$ ).

Taken as a whole, these results lead us to two conclusions. First, "corner" locations do not appear to be special in a reorientation context; subjects trained to visit wall locations produced behavior statistically indistinguishable from the more classic corner-visiting group. Second, that behavior suggests that subjects rely on features far more than on geometry in this particular reorientation paradigm, to the point where they almost fail to encode geometry altogether. This result is consistent with other reorientation research that has shown that, in some conditions, feature cues can overshadow geometric cues (e.g., Bodily, Eastman, & Sturz, 2011; Horne & Pearce, 2009; Pearce et al., 2004), but such results are not at all universal.

### Study 2: Eight locations

#### Method

**Subjects** Subjects were 35 University of Alberta undergraduates (22 female), who received course credit for participation. As in Study 1, participants were required to have normal color vision.

**Apparatus** The apparatus was identical to that employed in Study 1, except in regard to stimuli.

**Stimuli** The stimuli had the same general nature as in Study 1; only the possible color patterns were different. Four of the patterns were the same stimuli present in Study 1. The remaining four were brown vertical stripes, purple spots, black diagonal stripes (upper-right to lower-left), and orange hexagons. These stimuli are depicted in Fig. 3 and were placed in the configuration depicted in Fig. 1c.

As in Study 1, there were three transformed arenas as well. Transformation consisted of shuffling which color was present at a particular "slot." An affine transformation was created by shifting the colors two "slots" clockwise relative to training, which placed targets at novel wall configurations, but with the same angle (90° or 180°) relative to training. A separate partial transformation was created by shifting the colors one "slot" clockwise; in this condition, both the wall configuration and angle were different, as compared with training. Also as in Study 1, a featureless condition was included where all locations had identical pure-white colors in place of their original patterns.

**General procedures** The procedures followed were identical in all ways to those in Study 1, including the posttesting retention test and progression to Study 3 upon completion.

#### Results and discussion

Five subjects failed the posttesting retention test, resulting in 30 subjects (19 female) included in the analysis. No evidence of a gender effect (males,  $M = 2.0$ ,  $SD = 0.45$ ; females,  $M = 1.95$ ,  $SD = 0.78$ ;  $d' = 0.08$ , 95 % CI  $[-0.59, 0.94]$ ,  $p = .85$ ) or of

an effect of reinforcement grouping (group A,  $M = 2.1$ ,  $SD = 0.83$ ; group B,  $M = 1.8$ ,  $SD = 0.42$ ;  $d' = 0.51$ , 95 % CI  $[-0.27, 1.21]$ ,  $p = .28$ ) was found among the number of blocks these subjects required to complete training. Similarly, no significant difference was found in the number of blocks needed for subjects to learn that their reinforcement was at a corner or a wall (walls,  $M = 2.06$ ,  $SD = 0.44$ ; corners,  $M = 1.85$ ,  $SD = 0.86$ ;  $d' = 0.30$ , 95 % CI  $[0.33, 1.28]$ ,  $p = .45$ ).

In terms of choice consistency with particular cues, subjects exhibited patterns similar to those in Study 1 in the affine and featureless conditions, when adjusted to consider eight possible target locations. Within the affine arena, subjects in both the wall-target and corner-target conditions made similar proportions of choices consistent with geometry (wall,  $M = 0.26$ ,  $SD = 0.43$ ; corner,  $M = 0.33$ ,  $SD = 0.41$ ;  $d' = -0.17$ , 95 % CI  $[-0.89, 0.45]$ ,  $p = .60$ ) and consistent with features (wall,  $M = 0.74$ ,  $SD = 0.43$ ; corner,  $M = 0.61$ ,  $SD = 0.45$ ;  $d' = 0.31$ , 95 % CI  $[-0.36, 1.19]$ ,  $p = .376$ ). Similarly, the featureless arena saw subjects produce similar proportions of choices consistent with training geometry (wall,  $M = 0.49$ ,  $SD = 0.38$ ; corner,  $M = 0.54$ ,  $SD = 0.45$ ;  $d' = -0.12$ , 95 % CI  $[-0.88, 0.55]$ ,  $p = .71$ ). The partial condition, where both wall configuration and angle cues varied from training, also produced similar proportions of choices consistent with training geometry (wall,  $M = 0.30$ ,  $SD = 0.42$ ; corner,  $M = 0.31$ ,  $SD = 0.40$ ;  $d' = -0.03$ , 95 % CI  $[-0.68, 0.67]$ ,  $p = .90$ ) or with features (wall,  $M = 0.7$ ,  $SD = 0.42$ ; corner,  $M = 0.65$ ,  $SD = 0.42$ ;  $d' = 0.12$ , 95 % CI  $[-0.56, 0.87]$ ,  $p = .725$ ).

In every case, subjects from both groups produced indistinguishable results and were therefore pooled to test whether their choices varied from chance. Within the affine arena, subjects made choices consistent with the feature significantly more often than chance (95 % CI  $[0.54, 0.81]$ ,  $p < .05$ ) but did not make choices consistent with the original wall length configurations more often than chance (95 % CI  $[0.16, 0.43]$ ,  $p > .05$ ). Within the partial arena, subjects also followed the original wall length configuration at chance rates (95 % CI  $[0.17, 0.45]$ ,  $p > .05$ ) and followed feature cues significantly more often than chance (95 % CI  $[0.54, 0.80]$ ,  $p < .05$ ). Finally, unlike in Study 1, subjects responded to the original wall length configurations in the featureless arenas significantly more often than chance (95 % CI  $[0.39, 0.65]$ ,  $p < .05$ ).

In general, Study 2 supports Study 1's findings that walls are not significantly different from corners in terms of reorientation, even when the number of locations is increased. Features remain the best predictor of subject behavior and clearly dominate such behavior when they are presented in conflict with geometry, regardless of whether that geometry is completely inconsistent with training (the affine condition) or partially inconsistent with training (the partial condition). However, unlike in Study 1, subjects' behavior in the featureless arena clearly indicates that geometry was encoded during training. This is consistent with other literature that has shown that

humans are capable of completing reorientation tasks with more than four locations of interest (Newcombe et al., 2010).

This geometric overshadowing is the only noteworthy difference between Study 1's four-location task and Study 2's eight-location task. While this type of overshadowing has been observed during reorientation in some circumstances (Bodily et al., 2011; Pearce et al., 2004), this is usually not the case (e.g., Wall et al., 2004). It is quite possible that another version of the four-location reorientation task in which overshadowing does not occur may find no difference between four-location and eight-location reorientation, should it be repurposed for the latter. Testing this claim would require another experiment, beyond the scope of the present work.

### Study 3: Direct comparison

#### Method

**Subjects** Subjects were participants from both Study 1 and Study 2, as described above. This included subjects who failed their original study's posttesting retention test.

**General procedures** After completing the posttesting retention test for their original study, subjects found themselves placed back in the "welcome" chamber and were instructed that this marked the halfway point of the experiment. The protocol followed from here is identical to that described in Study 1, except for the stimulus set from whichever study the subject had not already seen. Following completion of this second experiment, subjects were debriefed.

#### Results and discussion

For discussions of the "first task," subjects who passed their posttesting retention tests as described in Study 1 and Study 2 were included in analysis. For the "second task," a separate posttesting retention test was performed (independently of the subjects' test result in their original study), and subjects who performed poorer than 70 % on this trial were excluded from analysis. Of the Study 1 subjects, 2 failed this test, leaving 34 (28 female), while 3 of the Study 2 subjects failed, leaving 34 (20 female) for analysis.

A large, and significant, order effect was observed (initial study,  $M = 1.94$ ,  $SD = 0.98$ ; second study,  $M = 1.36$ ,  $SD = 0.60$ ;  $d' = 0.7$ , 95 % CI  $[0.41, 1.05]$ ,  $p = .001$ ), regardless of which task was completed first. For the first task, there was no significant difference in training blocks between the four-location task and the eight-location task (four-location,  $M = 1.90$ ,  $SD = 1.25$ ; eight-location,  $M = 1.97$ ,  $SD = 0.67$ ;  $d' = -0.06$ , 95 % CI  $[-0.77, 0.37]$ ,  $p = .80$ ). There was also no evidence of any significant differences in time taken to learn the second task (four-location,  $M = 1.28$ ,  $SD = 0.46$ ; eight-location,  $M = 1.44$ ,  $SD = 0.70$ ;  $d' = -0.27$ , 95 % CI  $[-0.68, 0.18]$ ,  $p = .27$ ).

Two general conclusions can be drawn from these results. First, we do not find any evidence that changing the number of locations makes a task easier or harder to learn, both for naïve subjects and for subjects who have been trained in a different task. This conclusion is consistent with the operant perceptron's predictions but inconsistent with models that adjust for the total number of locations of interest (such as Miller & Shettleworth, 2007, 2008). Second, this result, combined with the order effect, allows us to conclude that subjects learned the second task faster than the first task, regardless of the number of locations present in either task. This suggests that something facilitated the second task. Since this is a difference in the mean number of training blocks, as opposed to a difference in actual time taken to complete those training blocks, it cannot be an increase in familiarity with the virtual world. Those familiar with the virtual world would be more proficient at positioning their virtual avatar where they wanted, but this would translate into a faster time per training block, not fewer total blocks required to learn the task.

One possibility is that, like the perceptron model discussed earlier, human subjects may be using a similar system to learn both tasks—a system in which the total number of locations is irrelevant. The perceptron accomplishes this by using an encoding that relies purely on local cues, where each location is considered independently of other locations during training. Alternative models based on matching current visual stimuli to previously learned visual stimuli (e.g., Cheung, Stürzl, Zeil, & Cheng, 2008; Wystrach, Cheng, Sosa, & Beugnon, 2011) have a similar property, while the associative models of Miller and Shettleworth (2007, 2008) do not, by virtue of their global probability terms.

## Model evaluation

With both simulation results and experimental evidence in analogous tasks at hand, we return now to the operant

**Table 3** Comparing human and network response frequencies with specific cue types

Task	Transform	Response Rates			
		Geometric		Feature	
		Human (95 % CI)	Network	Human (95 % CI)	Network
Four	Affine	[0.08,0.24]	0.49	[0.73, 0.90]	0.50
	Generalized	N/A	N/A	[0.82,0.95]	0.57
	Featureless	[0.46, 0.68]	0.83	N/A	N/A
Eight	Affine	[0.16, 0.43]	0.44	[0.54, 0.81]	0.45
	Partial	[0.17, 0.45]	0.51	[0.54, 0.80]	0.36
	Featureless	[0.39, 0.65]	0.71	N/A	N/A


perceptron and evaluate where the patterns of behavior agree and where they disagree. The operant perceptron made two broad classes of prediction: that subjects learn four-location reorientation and eight-location reorientation with equivalent amounts of training, and within each task, no appreciable difference exists between wall locations and corner locations in terms of reorientation behavior.

To see whether our experimental data agreed with the perceptron's claim of equivalent training time, we turned to a Bayesian analysis of results to generate likelihoods (and a corresponding semantic interpretation) for the null hypothesis, rather than leaving the results to a simple significance test (Gallistel, 2009). Within Study 1's four-location conditions, the odds in favor of a similar training time for walls and for corners are 2.88:1 (modest); within Study 2's eight-location conditions, these odds were 3.48:1 (substantial) in favor of no difference. For Study 3, for the first task, the hypothesis that the amount of time taken to learn to orient in four-location arenas and in eight-location arenas was the same had odds of 1.50:1 (weak); for the second task, the odds in favor of an equal amount of training time were 4.20:1 (substantial). The lower odds are not unexpected, given the relatively small number of subjects in these experiments. Importantly, in every case mentioned here, the odds are in favor of a null hypothesis of no difference in training time; an interesting and possibly counterintuitive prediction that arose from the computer simulation was supported by the experimental studies that used human subjects.

As for the second major prediction, of no significant behavioral differences between wall and corner locations, we turn to the pattern of perceptron response frequencies reported in the Appendix (here, these are the conditional probabilities of responding to a location given its cues, normalized across all locations within an arena). These tables also report the response frequencies for human subjects from Study 1 and Study 2. Are the networks' response frequencies appreciably different from the humans', or are they a plausible model?

To evaluate this, for each experimental condition, a 95 % confidence interval was bootstrapped onto the mean of the human subjects' responses to each type of cue (Efron & Tibshirani, 1994). These confidence intervals were used above to compare the response rates with chance; here, they are used to compare the response rate with the mean network response frequency to the same cue type (such frequencies are the normalized response rate). Since no difference was found in human subjects between wall groups and corner groups, the networks' corresponding location response frequencies were averaged as well. These confidence intervals and their comparisons are given in Table 3. In addition, Fig. 4 depicts the summarized human response frequencies for affine-transformed arenas from Study 1 and Study 2 (for information on the other conditions, refer to the Appendix); this can be qualitatively compared with Fig. 2, which reports responses to the same conditions in the perceptron.

**Fig. 4** Example human response frequencies in affine-transformed arenas. The figures are presented in the same manner as the perceptron data in Fig. 2. Refer to the Appendix for complete response data in all conditions



In all cases with four locations, the network model consistently predicted too many choices consistent with wall length geometry and too few choices consistent with features. This is not surprising, since humans consistently chose geometric-consistent locations in this task at a rate indistinguishable from chance, indicating that they did not encode geometric cues, while a review of the connection weights in Table 2 indicates that the networks did encode such cues. In the eight-location task, the networks performed much closer to human behavior in general, although again the networks tend to respond more frequently to geometry and less frequently to features than do humans. Qualitatively, this is supported by comparing Figs. 2 and 4: Except for humans trained in the four-location corner condition, rotational error is present in all locations to a similar degree, although the humans followed the feature to a greater extent.

These discrepancies in response frequencies suggest a need for exploring alternative design decisions in the perceptron. In particular, one open question concerns how changing the encoding of cue patterns might affect perceptron responses, as well as relationships between the networks and the humans that are based upon response frequency measures. Alternatively, the brightly colored features present during the experiment were highly visually salient, as compared with the matte-gray walls, and it is possible that the perceptron model (which assigned equal salience to all cue types) did not capture this difference. Exploring different learning rate parameters for each cue type, or conducting an experiment with less-salient features, may produce less of a discrepancy.

## General discussion

The present article explored reorientation in novel variations on the standard reorientation task, informed by a simple artificial neural network. Network simulations

and experimental data allowed us to examine how behavior changes—or rather, does not change—when the locations of interest are placed at locations other than the corners of a quadrilateral arena. Furthermore, we examined this in the context of changing the number of salient targets from four to eight and found that the difficulty in learning the task does not actually increase and, furthermore, that skills learned in one task generalize to the other in both directions. These results were consistent within both simulation and experiment and suggest that the behavior governing reorientation involves processing the cues available at locations taken in isolation, regardless of the global structure of the environment. In other words, learning reorientation does not require comparing the current location with any other possible location, contrary to the proposal of Miller and Shettleworth's (2007, 2008) model.

These results—both simulated and experimental—suggest an interesting refinement to the hypothesis that angles are processed in a manner similar to features (Sturz, Forloines, & Bodily, 2012). Under this refinement, the location of interest serves as a visually salient *focal point*—a reference from which wall length and angle cues are determined. In a typical reorientation task, the corners of an arena create that focal point (observe, for instance, how clear the boundaries are at the corners in Fig. 3b), but our results suggest that other visually salient goals (here, pillars, but also possibly bowls of food, boxes with toys, and so forth) can create the same effect, even if they are not placed at corners. The angle the walls form at the location of interest therefore becomes a feature of that location, in much the same manner as traditional feature cues, such as color. This allows associative models to capture associative strength at locations other than corners—a property demanded of these models by recent empirical work (Horne, León, & Pearce, 2013).

Furthermore, the negative weights present on certain cues and biases in the model suggest that inhibition is key to this process; that is, agents not only learn to approach cues that were

reinforced during training, but also learn to *avoid* cues that were *not* reinforced during training. Classical reorientation results, such as rotational errors, may not emerge through a process based on relative attractiveness (e.g., Miller & Shettleworth, 2008) but may, rather, emerge due to the interplay between excitation and inhibition when presented with a transformed arena. In particular, the negative biases in the model suggest that the model is inherently hesitant about investigating locations unless certain cues are observed. When presented with a dramatically transformed arena (i.e., the four-location generalized condition, where networks trained in corners are tested at walls, and vice versa), responses resembling rotational error emerge due to the *absence of such inhibitory cues*, rather than the presence of excitatory ones. This suggests that investigating the role of inhibition may play a critical role in the development of our theoretical understanding of spatial learning.

Additionally, theories of viewpoint-matching (Cheung et al., 2008) propose that reorientation is largely a matter of learning broad visual stimuli when reinforced, then seeking to minimize the difference between one's current visual input and this learned image. This theory could also be capable of reorienting in arenas without corner-based locations. Interestingly, both this theory and the operant perceptron model learning as error-correcting on the basis of a pattern of subject behavior, although they encode the available stimuli in dramatically different manners. It would be interesting to see how their predictions on choice frequencies differ, if at all. Evaluations and comparisons of earlier associative and viewpoint-matching theories exist in the context of ant navigation (Wystrach et al., 2011), which suggest that associative models' inherent segregation of feature and geometry cues is problematic. Although the encoding presented to the operant perceptron does code features separate from "geometry," it also suggests that the two are processed in the same manner (as the presence or absence of particular cue indicators).

The choice to use configuration unit encoding (i.e., a different unit turns on for each possible wall-length configuration) choice was made for consistency with existing literature (Dupuis & Dawson, *in press*; Horne & Pearce, 2010; Miller & Shettleworth, 2007, 2008). If the simulation diverges from live-agent data, then this theoretical choice may not be appropriate. Indeed, a divergence appears in two aspects: The simulation encoded geometry in the four-location task, while human subjects did not, and the networks produced slightly different choice frequencies (Table 3). However, in spite of this, the model still correctly predicted several interesting results, such as the lack of a difference between corner-trained and wall-trained subjects and the similar difficulty of the four-location and eight-location tasks. This suggests that the model needs adjustment, but that adjustment need not be extreme. This adjustment could take the form of parametric adjustment (i.e., salience changes, discussed above) or a change in how

arena walls are encoded (i.e., thermometer coding; Dawson, Kelly, et al., 2010); these changes may have important theoretical implications, if any prove more fruitful than configuration encoding.

These developments provide examples of experimental results informing future modeling decisions, which is a common practice in cognitive science. However, the present work, in contrast, demonstrated that modeling can quite easily inform experiment as well. A new experimental result or theoretical construct can revise an existing model, which in turn can be used to generate empirical claims in novel environments quickly and cheaply. If any of those predictions are of interest, future experimentation can be used to test these new hypotheses.

This methodological style—creating simple and plausible models that *behave* and then generating hypotheses and experiments on the basis of this behavior—is an example of the synthetic approach to cognitive science (Dawson, 2004; Dawson, Dupuis, & Wilson, 2010). The synthetic approach can prove fruitful in breaking future deadlock or opening up novel research paradigms. For example, to the best of our knowledge, the present work describes the first attempt at systematically varying the nature and number of locations during reorientation. The decision to investigate this comparison was motivated entirely by the structure of the neural network model. Additionally, although the interpretation of angle as a feature cue is contested in the literature (Hupbach & Nadel, 2005; Lee & Spelke, 2011) and the traditional response would be to design an experiment to fit data to see whether such an interpretation is valid, here, the synthetic approach allows that problem to be circumvented to a point. We posit a model that processes angles as features and then see what it can and cannot do. Here, that model behaved in a way consistent with our human subjects, with their largest differences being due to parametric selection rather than a structural change. Our neural network model suggests that processing angles as features may be entirely consistent with associative theories of reorientation.

This neural network model is a ripe avenue for future research. The operant perceptron successfully handled reorientation with assorted numbers of target locations, positioned at arbitrary points along the edge of an arena, using the encoding described above. This encoding can easily be extended, including extra input units to represent angles other than 90° or 180° or to represent other sets of wall length configurations than those used here. Putting these properties together allows this architecture to handle any polygonal arena with edge-defined locations of interest, including kites (Dawson, Kelly, et al., 2010; Pearce et al., 2004), hexagons (Sturz & Bodily, 2011), and octagons (Newcombe et al., 2010). While we could extend the operant perceptron to see whether it fits the data from some of

the novel tasks (i.e., regular octagons; Newcombe et al., 2010) in a manner similar to a more standard perceptron (Dawson, Kelly, et al., 2010), the synthetic approach would be to generate totally new predictions inspired by our findings. In this case, we might try nonuniform octagons (an arena type not yet investigated), or we might note other successes of the operant perceptron altogether, such as superconditioning (Dupuis & Dawson, *in press*) or probability matching (Dawson et al., 2009), and branch out beyond reorientation into completely new paradigms.

Lewandowsky (1993) observed that computer modeling had its benefits, if done with care. The present work illustrates all of these core ideas. A desire to increase mathematical rigor in the Miller and Shettleworth (2007, 2008) model led to the development of new tools—both the operant perceptron

(Dupuis & Dawson, *in press*) used in simulation and the fAARS-Lite platform (Dupuis, 2012; Gutiérrez, 2012; Lubyk et al., 2012) used in data collection. These tools facilitated finding and testing the tacit assumption in reorientation literature that corners have some inherently special property. Finally, the simulation results indicating that there should be no difference in effort needed to learn reorientation in arenas with more locations fit the description for serendipitous findings in novel environments, a point emphasized by the same result appearing among human subjects. It would appear that, even after 20 years, Lewandowsky's observations and advice for cognitive modelers still remains effective.

**Author Note** This work is based on part of a master's thesis defended in September 2012 for the Department of Psychology, University of Alberta.

## Appendix

**Table 4** Average responses of operant perceptrons and human subjects to the four-location simulation

Task	Type	Transform	Loc	Color	Angle	Wall Configuration		Network		Human Frequency
						Left	Right	Activity	Frequency	
Four	Corner	Affine	1	Blue	90*	6*	3*	0.36	0.24	0.06
			2	Green*	90*	3	6	0.75	0.50	0.89
			3	Red	90*	6*	3*	0.36	0.24	0.03
			4	Yellow	90*	3	6	0.01	0.01	0.02
		Generalized	1	Blue	180	2	4	0.26	0.17	0.04
			2	Green*	180	2	1	0.94	0.62	0.87
			3	Red	180	2	4	0.26	0.17	0.04
			4	Yellow	180	2	1	0.06	0.04	0.04
	Featureless	1	White	90*	6*	3*	0.52	0.42	0.26	
		2	White	90*	3	6	0.11	0.08	0.19	
		3	White	90*	6*	3*	0.52	0.42	0.30	
		4	White	90*	3	6	0.11	0.08	0.25	
	Wall	Affine	1	Yellow	180*	2*	4*	0.37	0.25	0.16
			2	Blue*	180*	2	1	0.75	0.50	0.68
			3	Green	180*	2*	4*	0.36	0.24	0.11
			4	Red	180*	2	1	0.01	0.01	0.05
		Generalized	1	Green	90	6	3	0.26	0.22	0.02
			2	Red	90	3	6	0.06	0.05	0.05
			3	Yellow	90	6	3	0.26	0.22	0.03
			4	Blue*	90	3	6	0.60	0.51	0.49
Featureless		1	White	180*	2*	4*	0.53	0.42	0.26	
		2	White	180*	2	1	0.11	0.08	0.22	
		3	White	180*	2*	4*	0.53	0.42	0.35	
		4	White	180*	2	1	0.11	0.08	0.29	

\* Cue type reinforced during training



**Table 5** Average responses of operant perceptrons and human subjects to the eight-location simulation

Task	Type	Transform	Loc.	Color	Angle	Wall Configuration		Network		Human	
						Left	Right	Activity	Frequency	Frequency	
Eight	Corner	Affine	1	Black	180	2	4	0.04	0.03	0.01	
			2	Blue	90*	4*	2*	0.36	0.22	0.16	
			3	Purple	180	2	1	0.04	0.03	0.02	
			4	Green*	90*	1	2	0.74	0.45	0.58	
			5	Orange	180	2	4	0.04	0.03	0.00	
			6	Red	90*	4*	2*	0.36	0.22	0.18	
			7	Brown	180	2	1	0.04	0.03	0.04	
			8	Yellow	90*	1	2	0.01	0.01	0.01	
		Partial	1	Blue	180	2	4	0.03	0.02	0.01	
			2	Purple	90*	4*	2*	0.42	0.25	0.19	
			3	Green*	180	2	1	0.61	0.36	0.65	
			4	Orange	90*	1	2	0.07	0.04	0.00	
			5	Red	180	2	4	0.03	0.02	0.00	
			6	Brown	90*	4*	2*	0.42	0.25	0.11	
			7	Yellow	180	2	1	0.01	0.00	0.03	
			8	Black	90*	1	2	0.07	0.04	0.01	
		Featureless	1	White	180	2	4	0.05	0.04	0.00	
			2	White	90*	4*	2*	0.49	0.35	0.26	
			3	White	180	2	1	0.05	0.04	0.06	
			4	White	90*	1	2	0.09	0.07	0.15	
			5	White	180	2	4	0.05	0.04	0.01	
			6	White	90*	4*	2*	0.49	0.35	0.26	
			7	White	180	2	1	0.05	0.04	0.09	
			8	White	90*	1	2	0.09	0.07	0.17	
	Wall	Affine	1	Black	180*	2*	4*	0.36	0.22	0.11	
			2	Blue	90	4	2	0.04	0.03	0.00	
			3	Purple*	180*	2	1	0.74	0.45	0.76	
			4	Green	90	1	2	0.04	0.03	0.00	
			5	Orange	180*	2*	4*	0.36	0.22	0.13	
			6	Red	90	4	2	0.04	0.03	0.00	
			7	Brown	180*	2	1	0.01	0.01	0.00	
			8	Yellow	90	1	2	0.04	0.03	0.00	
			Partial	1	Blue	180*	2*	4*	0.42	0.25	0.18
				2	Purple*	90	4	2	0.61	0.36	0.71
				3	Green	180*	2	1	0.07	0.04	0.00
				4	Orange	90	1	2	0.03	0.02	0.00
				5	Red	180*	2*	4*	0.43	0.25	0.11
				6	Brown	90	4	2	0.01	0.00	0.00
				7	Yellow	180*	2	1	0.07	0.04	0.00
				8	Black	90	1	2	0.03	0.02	0.00
Featureless		1	White	180*	2*	4*	0.50	0.35	0.29		
		2	White	90	4	2	0.05	0.04	0.15		
		3	White	180*	2	1	0.09	0.07	0.13		
		4	White	90	1	2	0.05	0.04	0.00		
		5	White	180*	2*	4*	0.50	0.35	0.18		
		6	White	90	4	2	0.05	0.04	0.14		
		7	White	180*	2	1	0.09	0.07	0.07		
		8	White	90	1	2	0.05	0.04	0.03		

\* Cue type reinforced during training

## References

- Bodily, K. D., Eastman, C. K., & Sturz, B. R. (2011). Neither by global nor local cues alone: Evidence for a unified orientation process. *Animal cognition*, *14*, 665–674. doi:10.1007/s10071-011-0401-x
- Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition*, *23*(2), 149–78. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/3742991>
- Cheng, K. (2008). Whither geometry? Troubles of the geometric module. *Trends in cognitive sciences*, *12*(9), 355–61. doi:10.1016/j.tics.2008.06.004
- Cheng, K., & Gallistel, C. R. (2005). Shape parameters explain data from spatial transformations: Comment on Pearce et al. (2004) and Tommasi & Polli (2004). *Journal of experimental psychology. Animal behavior processes*, *31*(2), 254–9. doi:10.1037/0097-7403.31.2.254. discussion 260–1.
- Cheng, K., & Newcombe, N. S. (2005). Is there a geometric module for spatial orientation? Squaring theory and evidence. *Psychonomic bulletin & review*, *12*(1), 1–23. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15945200>
- Cheung, A., Stürzl, 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*(1). doi:10.1037/0097-7403.34.1.15
- Dawson, M. R. W. (2004). *Minds and Machines: Connectionism and Psychological Modeling*. Oxford, UK: Blackwell.
- Dawson, M. R. W. (2008). Connectionism and Classical Conditioning. *Comparative Cognition and Behaviour Reviews*, *3*, 115. Comparative Cognition Society. doi:10.3819/ccbr.2008.30008
- Dawson, M. R. W., Dupuis, B., Spetch, M. L., & Kelly, D. M. (2009). Simple artificial neural networks that match probability and exploit and explore when confronting a multiarmed bandit. *IEEE Transactions in Neural Networks*, *20*(8), 1368–1371. doi:10.1109/TNN.2009.2025588
- Dawson, M. R. W., Dupuis, B., & Wilson, M. (2010a). *From Bricks to Brains: The Embodied Cognitive Science of LEGO Robots* (p. 331). Edmonton, AB: Athabasca University Press.
- 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*(3), 415–8. doi:10.1037/0097-7403.34.3.415
- Dawson, M. R. W., Kelly, D. M., Spetch, M. L., & Dupuis, B. (2010b). Using perceptrons to explore the reorientation task. *Cognition*, *114*(2), 207–26. doi:10.1016/j.cognition.2009.09.006
- Dupuis, B. (2012). *The Cognitive Science of Reorientation*. (Master's thesis.) University of Alberta. Retrieved from <http://hdl.handle.net/10402/era.29120>
- Dupuis, B., & Dawson, M. R. W. (in press). Differentiating Models of Associative Learning: Reorientation, Superconditioning, and the Role of Inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*.
- Edgington, E. S. (1995). *Randomization tests* (3rd ed.). New York, NY: Marcel Dekker, Inc.
- Efron, B., & Tibshirani, R. J. (1994). *Introduction to the bootstrap*. New York, NY: Chapman & Hall.
- Gallistel, C. R. (1990). *The Organization of Learning*. Cambridge, MA: MIT Press.
- Gallistel, C. R. (2009). The importance of proving the null. *Psychological review*, *116*(2), 439–453. doi:10.1037/a0015251.The
- Gutiérrez, L. A. (2012). *The fAARS Platform For Augmented Alternate Reality Services and Games*. University of Alberta. Retrieved from <https://era.library.ualberta.ca/public/view/item/uuid:b8d36cef-585d-4010-8124-9a22c1384087>
- Horne, M. R., León, S. P., & Pearce, J. M. (2013). The influence of excitatory and inhibitory landmarks on choice in environments with a distinctive shape. *Journal of experimental psychology. Animal behavior processes*, *39*(1), 76–84. doi:10.1037/a0030459
- Horne, M. R., & Pearce, J. M. (2009). Between-cue associations influence searching for a hidden goal in an environment with a distinctive shape. *Journal of Experimental Psychology: Animal Behavior Processes*, *35*(1), 99–107. doi:10.1037/0097-7403.35.1.99
- Horne, M. R., & Pearce, J. M. (2010). Conditioned inhibition and superconditioning in an environment with a distinctive shape. *Journal of Experimental Psychology: Animal Behavior Processes*, *36*(3), 381–94. doi:10.1037/a0017837
- Hupbach, A., & Nadel, L. (2005). Reorientation in a rhombic environment: No evidence for an encapsulated geometric module. *Cognitive Development*, *20*(2), 279–302. doi:10.1016/j.cogdev.2005.04.003
- Huttenlocher, J. (2003). How toddlers represent enclosed spaces. *Cognitive Science*, *27*(5), 749–766. doi:10.1016/S0364-0213(03)00062-4
- Lee, S. A., & Spelke, E. S. (2011). Young children reorient by computing layout geometry, not by matching images of the environment. *Psychonomic bulletin & review*, *18*(1), 192–8. doi:10.3758/s13423-010-0035-z
- Lewandowsky, S. (1993). The rewards and hazards of computer simulations. *Psychological science*, *4*(4), 236–243.
- Lubyk, D. M., Dupuis, B., Gutiérrez, L. A., & Spetch, M. L. (2012). Geometric orientation by humans: Angles weigh in. *Psychonomic Bulletin & Review*, *19*(3), 436–442. doi:10.3758/s13423-012-0232-z
- Lubyk, D. M., & Spetch, M. L. (2012). Finding the best angle: Pigeons (*Columba livia*) weight angular information more heavily than relative wall length in an open-field geometry task. *Animal cognition*, *15*(3), 305–12. doi:10.1007/s10071-011-0454-x
- Miller, N. Y. (2009). Modeling the effects of enclosure size on geometry learning. *Behavioural processes*, *80*(3), 306–13. doi:10.1016/j.beproc.2008.12.011
- Miller, N. Y., & Shettleworth, S. J. (2007). Learning about environmental geometry: An associative model. *Journal of Experimental Psychology: Animal Behavior Processes*, *33*(3), 191–212. doi:10.1037/0097-7403.33.3.191
- Miller, N. Y., & Shettleworth, S. J. (2008). An associative model of geometry learning: A modified choice rule. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*(3), 419–22. doi:10.1037/0097-7403.34.3.419
- Newcombe, N. S., Ratliff, K. R., Shallcross, W. L., & Twyman, A. D. (2010). Young children's use of features to reorient is more than just associative: Further evidence against a modular view of spatial processing. *Developmental science*, *13*(1), 213–20. doi:10.1111/j.1467-7687.2009.00877.x
- Pearce, J. M., Good, M. A., Jones, P. M., & McGregor, A. (2004). Transfer of spatial behavior between different environments: Implications for theories of spatial learning and for the role of the hippocampus in spatial learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *30*(2), 135–47. doi:10.1037/0097-7403.30.2.135
- Ratliff, K. R., & Newcombe, N. S. (2008). Reorienting when cues conflict: Evidence for an adaptive-combination view. *Psychological science*, *19*(12), 1301–7. doi:10.1111/j.1467-9280.2008.02239.x
- Rescorla, R. A. (1967). Pavlovian conditioning and its proper control procedures. *Psychological review*, *74*(1), 71–80. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/5341445>
- 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* (pp. 64–99). New York, NY: Appleton-Century-Crofts.
- Rosenblatt, F. (1962). *Principles of Neurodynamics*. Washington, DC: Spartan Books.

- Rummelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*, 533–536.
- Sturz, B. R., & Bodily, K. D. (2011). Is surface-based orientation influenced by a proportional relationship of shape parameters? *Psychonomic bulletin & review*, *18*(5), 848–854. doi:10.3758/s13423-011-0111-z
- Sturz, B. R., Forloines, M. R., & Bodily, K. D. (2012). Enclosure size and the use of local and global geometric cues for reorientation. *Psychonomic bulletin & review*, *19*(2), 270–6. doi:10.3758/s13423-011-0195-5
- Tommasi, L., & Polli, C. (2004). Representation of two geometric features of the environment in the domestic chick ( *Gallus gallus*). *Animal cognition*, *7*(1), 53–9. doi:10.1007/s10071-003-0182-y
- Twyman, A. D., & Newcombe, N. S. (2010). Five reasons to doubt the existence of a geometric module. *Cognitive science*, *34*(7), 1315–56. doi:10.1111/j.1551-6709.2009.01081.x
- 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*(3), 289–98. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15672824>
- Webb, B. (2000). What does robotics offer animal behaviour? *Animal behaviour*, *60*(5), 545–558. doi:10.1006/anbe.2000.1514
- Wystrach, A., Cheng, K., Sosa, S., & Beugnon, G. (2011). Geometry, features, and panoramic views: Ants in rectangular arenas. *Journal of experimental psychology. Animal behavior processes*, *37*(4), 420–35. doi:10.1037/a0023886