Cognition xxx (2009) xxx–xxx

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Using perceptrons to explore the reorientation task

Michael R.W. Dawson [a,*], Debbie M. Kelly [b], Marcia L. Spetch [a], Brian Dupuis [a]

[a] *Department of Psychology, University of Alberta, Edmonton, Alberta, Canada T6G 2P9*
[b] *Department of Psychology, University of Saskatchewan, Saskatoon, Saskatchewan, Canada S7N 5A5*

ARTICLE INFO

Article history:
Received 3 June 2009
Revised 2 September 2009
Accepted 5 September 2009
Available online xxxx

Keywords:
Reorientation task
Geometric cues
Feature cues
Artificial neural networks

ABSTRACT

The reorientation task is a paradigm that has been used extensively to study the types of information used by humans and animals to navigate in their environment. In this task, subjects are reinforced for going to a particular location in an arena that is typically rectangular in shape. The subject then has to find that location again after being disoriented, and possibly after changes have been made to the arena. This task is used to determine the geometric and featural cues that can be used to reorient the agent in the arena. The purpose of the present paper is to present several simulation results that show that a simple neural network, a perceptron, can be used to generate many of the traditional findings that have been obtained using the reorientation task. These results suggest that reorientation task regularities can be explained without appealing to a geometric module that is a component of spatial processing.

© 2009 Elsevier B.V. All rights reserved.

1. The reorientation task

The ability to orient and navigate in space is critical for the survival of humans and animals. Studies of navigation in indoor environments have found that humans and other animals can use available external cues to determine direction (Cheng & Newcombe, 2005). Such cues can include the overall shape of the environment (geometric cues), as well as other available landmarks or local elements that might also be placed in the environment (feature cues). Geometric cues are presumed to be relational, while feature cues are not: "A geometric property of a surface, line, or point is a property it possesses by virtue of its position relative to other surfaces, lines, and points within the same space. A non-geometric property is any property that cannot be described by relative position alone" (Gallistel, 1990, p. 212). One question of considerable interest is the extent to which either geometric or feature cues are used to govern navigation.

One approach that has been used extensively to answer this question is the reorientation task, first introduced by Cheng (1986). In this paradigm, an agent is placed within an enclosure or arena that is usually rectangular in shape. The metric properties of the arena (i.e., length of walls, angles between walls) combined with the distinction between left and right (e.g., the long wall is to the left of the short wall) provide geometric cues. Colors of walls, or the visual properties of additional objects added to the arena (e.g., placed at each corner of a rectangular enclosure) can be used to provide feature cues (see Fig. 1). In the reorientation task, an agent learns that a particular place – usually a corner of a rectangular arena – is a goal location. The agent is then removed from the arena, disoriented, and returned to an arena, with the task of using the available cues to relocate the goal. The agent can, of course, be placed back into the original, unaltered arena. Of more interest are experimental conditions in which the arena has been changed in some way.

For example, after training in one arena (e.g., Fig. 1B or D) the subject might be placed back into an arena after the feature cues have been moved to different locations (e.g., Fig. 1C or E). This manipulation places feature cues in conflict with geometric cues. Will the agent move to a location defined by geometric information, or will it move to a different location indicated by feature information? Extensive

---

* Corresponding author.
E-mail address: mdawson@ualberta.ca (M.R.W. Dawson).

0010-0277/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.cognition.2009.09.006

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

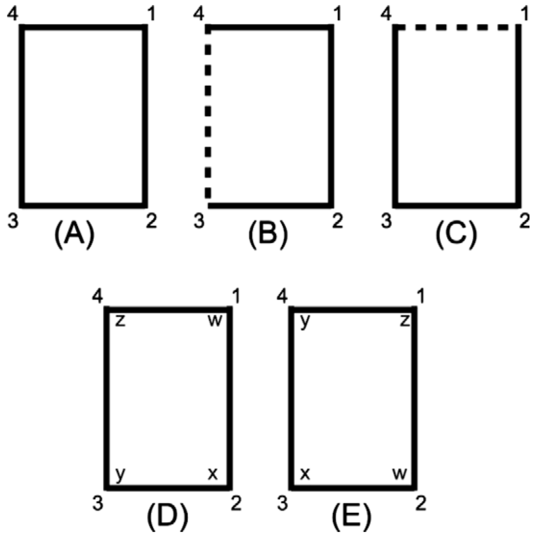

Fig. 1. Examples of rectangular arenas that can be used to study spatial reorientation. (A) A rectangular arena with no feature cues. The corners (Locations 1 through 4) are potential locations for reinforcement. Note that the corners at Locations 4 and 2 are geometrically equivalent to one another, as are the corners at Locations 1 and 3. (B) Wall color used as a feature cue. The wall indicated by the dashed line would be one color (e.g., white) while the other three walls would be a different color (e.g., black). (C) An affine transformation of (B), usually described as a conflict test when an animal is trained in an arena like that in (B), and then placed in this arena. In this conflict test, Locations 4 and 2 have correct geometry, but incorrect features. Location 1 has correct features, but incorrect geometry. Location 3 has incorrect geometry and incorrect features. (D) Feature cues as landmarks at each location. Each letter stands for a unique object (e.g., a colored or patterned panel) that can be used to identify the location. (E) An affine transformation of (D).

use of the reorientation task has uncovered a wealth of empirical evidence about animal navigation and the relative relevance of geometric and feature cues. These results have revealed some striking regularities in the reorientation task.

First, consider the case in which animals must perform the reorientation task in a rectangular arena using only geometric cues (Fig. 1A). This occurs when no unique feature cues are present in the arena. One consequence of this condition is that geometric cues do not specify a single target location in such an arena. For instance, the geometric cues available at Location 4 of Fig. 1A are identical to those available at Location 2 of the same figure: 90° angle, longer wall to the left and shorter wall to the right. As a result, these two corners are geometrically indistinguishable. When agents are trained on the reorientation task under such conditions, one of the basic findings is rotational error (Cheng, 1986, 2005). When rotational error occurs, the trained animal goes to the reinforced corner, as well as the corner located at a 180° rotation through the center of the arena, at above chance levels. That is, the agent cannot, and should not be able to, distinguish the reinforced corner from another corner that has identical geometric properties. This is usually taken as evidence that the animal is relying upon the geometric properties of the environment. Rotational error has been found in numerous studies with species ranging from ants (Wystrach & Beugnon, 2009) to humans (see Cheng and Newcombe (2005) for a review).

The second main regularity that governs the reorientation task occurs when feature cues (e.g., distinct objects) are added to the arena. These cues can be used by agents to uniquely relocate the reinforced location. For instance, feature cues can be added by making one of the arena walls a distinctive color (Fig. 1B), or by placing a unique landmark at each corner of the arena (Fig. 1D). The addition of such information can eliminate the response indeterminacy that is observed when only geometric cues are available.

Third, even though unique objects may be sufficient to correctly relocate a reinforced place in the arena, it would appear that in most cases agents use both feature and geometric cues. That is, geometric cues can influence behavior even when such cues are not required to solve the task. This claim is supported by several pieces of evidence. First, in some cases subjects continue to make some rotational errors even when a feature disambiguates the correct corner (Cheng, 1986; Hermer & Spelke, 1994). Second, when features are removed following training, subjects typically revert to choosing both of the geometrically correct locations (Kelly, Spetch, & Heth, 1998; Sovrano, Bisazza, & Vallortigara, 2003). Third, when the features are moved after training so as to create a conflict between geometric and feature cues, control by both types of information is often observed (Brown, Spetch, & Hurd, 2007; Kelly et al., 1998; Ratliff & Newcombe, 2008); the extent of control by geometric information on such tests appears to depend on several factors, including species, prior experience, and size of arena (Cheng & Newcombe, 2005). Thus, even when feature cues provide the most reliable indicator of the goal location, geometric information is typically also encoded.

Early theories of the regularities governing the reorientation task proposed that geometric features were encoded by modular processes that were dedicated to this kind of information (Cheng, 1986; Gallistel, 1990). For example, Gallistel (1990) viewed the solution of the reorientation task as a two stage process. The first stage occurs when an agent is first placed in an arena: it encodes the shape of the arena by attending to metric cues, such as wall lengths and angles between walls, as well as to sense cues (i.e., the distinction between left and right). The purpose of encoding the arena's shape is that this information is then used by the agent to determine its heading: that is, the arena's shape provides the reference frame for the agent's ability to orient itself. The second stage occurs when an agent is disoriented, and then placed in an arena once again. In this stage, the agent uses a representation of the shape of the previously encountered arena as a mental map. The agent "gets its heading and position on its map by finding the rotation and translation required to produce a congruence (shape match) between the currently perceived shape of the environment and a corresponding region of its map" (p. 220). If the only sources of information used to create such maps are sense and geometric cues, one consequence of this theory is rotational error in rectangular arenas.

A key assumption of the Gallistel (1990) model is that the processing of environmental shape is modular (Fodor,

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

1983). According to Fodor, a module is a neural substrate that is specialized for solving a particular information processing problem. A module has access only to limited information to use to solve this problem; other information is simply not available (even if it is relevant). Modules operate in a fast, mandatory fashion; they exhibit characteristic breakdown patterns when they fail because of their specialized neural circuitry; and they operate independently of the influence of the contents of higher-order beliefs – that is, they are cognitively impenetrable (Pylyshyn, 1984). It has been argued (Cheng, 1986; Gallistel, 1990) that the geometric computations in Gallistel's model are modular because they are mandatory (i.e., they are carried out even when unnecessary because of the presence of feature cues) and impenetrable because they are not influenced by "information about surfaces other than their relative positions" (Gallistel, 1990, p. 208), which would be useful to disambiguate the correct and rotational corners.

Recently, questions have been raised about this strong assumption that geometric cues are processed by a geometric module. One reason for this is that the most compelling evidence for claims of modularity comes from neuroscience (Dawson, 1998; Fodor, 1983), but such evidence about the modularity of geometry in the reorientation task is admittedly sparse (Cheng & Newcombe, 2005). As a result, most arguments about modularity in this context are based on behavioral data. However, such data is consistent with many different notions of modularity (e.g., Cheng & Newcombe, 2005, Fig. 3.

One result of the questions that have been raised about the existence of a geometric module is the proposal of alternative notions of modularity (e.g., Cheng, 2005; Cheng & Newcombe, 2005). Cheng (2005, p. 17), suggests that "geometric and feature information are encoded together in one record for localization. This process is non-modular." Cheng then proceeds to preserve modularity by arguing that different types of information might be stored in the same location, but when certain devices access this common store, they only access particular types of information, and are thus modular in nature. In short, Cheng conjoins "a modular process and a non-modular representational structure." Interestingly, this is exactly the sort of processing that typifies nonmodular information processing architectures, such as production systems (Newell, 1973, 1980), where each processor searches a working memory for very particular triggering information.

More recently, some researchers have considered abandoning the assumption of geometric modularity completely (Cheng, 2008). The first alternative to modularity was adaptive combination theory (Newcombe & Ratliff, 2007). In this theory, mechanisms are used to weight contributions of different sources of information when they are combined to control a decision or judgment. The weighting can combine information sources that would not all be available to a single module, and reflects an agent's history of experience with this information as well as its variability (for instance, when the weighting is governed by Bayesian theory (Cheng, Shettleworth, Huttenlocher, & Rieser, 2007). Another recent model of the reorientation task uses a general theory of operant associative learning in which geometric and feature cues are not treated differentially (Miller & Shettleworth, 2007, 2008). A third theory predicts reorientation task behavior by arguing that the agent maximizes the visual similarity (on the basis of very raw images) of locations in the new arena to the image of the goal location in the original arena (Cheung, Stuerzl, Zeil, & Cheng, 2008; Stuerzl, Cheung, Cheng, & Zeil, 2008). In this theory, the metric of visual similarity does not make explicit the geometric properties that were central to original theories of the task. Another theory attempts to explain reorientation task regularities using a robot that reacts to environmental stimuli, and does not model the arena at all (Nolfi, 2002; Nolfi & Floreano, 2000).

As researchers revisit, revise, and potentially abandon the notion of geometric modularity (Cheng, 2005, 2008; Cheng & Newcombe, 2005), it is timely to consider alternative approaches. How might this be accomplished? The possibility explored below is to adopt the paradigm of synthetic psychology (Braitenberg, 1984; Brooks, 1999; Clark, 1997, 2003; Dawson, 2004; Pfeifer & Scheier, 1999): that is, to build extremely simple models, and to use these models to explore the extent to which they can produce interesting behavior when interacting with their environment. In particular, we explored the ability of a very simple artificial neural network to deal with a number of different versions of the reorientation task. This network does not include a "geometric module", and does not qualitatively differentiate geometric from featural information. If such networks can generate a number of regularities that govern the reorientation task, then they provide a plausible, and non-modular, theory of reorientation.

In the remainder of this paper we briefly introduce artificial neural networks, and describe in detail how the reorientation task can be defined for a particular type of network, the perceptron. We then report the results of a number of simulations that explore the ability of this type of network to provide insight about the reorientation task.

2. Perceptrons and the reorientation task

An artificial neural network is a system of simple, interconnected processing units that learns to generate a desired response to a presented stimulus by adjusting the weights of the connections between its processors (Bechtel & Abrahamsen, 2002; Dawson, 2004, 2005; Rumelhart & McClelland, 1986). Typically, artificial neural networks are composed of three categories of processors: input units, hidden units, and output units. The stimulus is encoded as a pattern of activity in a set of input units. The response is represented as a pattern of activity in a set of output units. The hidden units are intermediate processors that detect complex features in the stimulus, which in turn determine the network's response. A simpler, and less powerful, type of artificial neural network is a perceptron (Rosenblatt, 1958, 1962), which does not include a layer of hidden units. While this type of network is simpler than more modern multilayered architectures, it still has a great deal of psychological relevance. Perceptrons can be used to simulate a great many results in the classical conditioning

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

literature (Dawson, 2008), and the training procedure for such networks is formally equivalent to the Rescorla–Wagner model of animal learning (Gluck & Bower, 1988; Sutton & Barto, 1981). Indeed, the simulations below show that perceptrons are capable of solving the reorientation task. We now describe in more detail how the reorientation task can be represented for such networks.

The perceptrons involved in the simulations below all used input units to encode various stimulus features present at different locations in a reorientation task arena. They also had a single output unit used to encode a response to each stimulus, which could be interpreted as the perceptron's judgment of the likelihood of receiving a reward at a stimulus location. Each input unit was connected to the output unit by a weighted connection whose value was determined by a process of associative learning. An illustration of the perceptrons used in the simulations below is provided in Fig. 2.

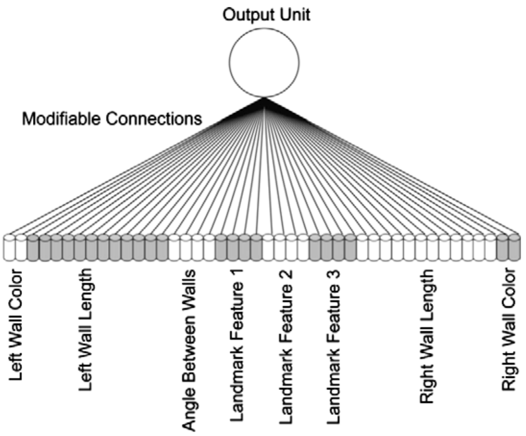

Fig. 2. The perceptrons used in the simulations had a single output unit, and 44 input units. The output unit used a logistic activation function. The input units are represented by the cylinders, and the lines indicate modifiable connections between input units and the output unit. The figure uses shading to indicate groups of input units that are used to code a particular cue; the cue encoded by a group is labeled in the figure. The coding scheme used for each cue type is described in the Table 1 in the text. In this coding scheme, input units were turned either on or off to represent various feature values at a location of interest.

2.1. *Defining the reorientation task for a perceptron: stimuli*

Consider a perceptron to be an agent that inspects each of the four corners of the rectangular arena illustrated in Fig. 1. For each inspected location, the perceptron registers a number of different properties: the length and color of the wall to the left, the length and color of the wall to the right, the angle between the walls at the location, and three different features (e.g., color, shape, height) of an object that might be present. The different property values were encoded by activating the perceptron's input units. Forty-four different input units were required to represent the eight different features that could be available at any location. Table 1 provides the list of features, the units involved in encoding a feature, and examples of encodings of different possible values of these features. In some simulations, one or more features may not be present. The absence of a feature was represented by activating all of its associated input units with a value of 0 (i.e., turning all of the feature's units off).

Note that the encoding scheme that was employed used local codes for some features (i.e., all those represented by unary coding). In a local code, turning a single input unit on provides enough information to define a feature (e.g., Dawson, 2008, Chapter 9). The encoding scheme for wall length was distributed. In a distributed code, the definition of a feature requires the simultaneous consideration of signals from multiple input units, as is the case in a thermometer code for distance. The distinction between local and distributed codes is different from the distinction between elemental and configural representations that is important to consider in neural network models of biological learning (Pearce, 1997, 2002). Both elemental and configural representations can be implemented in a perceptron using local encodings (e.g., Dawson, 2008, Section 9.6). An elemental representation is accomplished by encoding the presence

Table 1
Encoding of a location's properties using 44 input units of a perceptron.

Input units	Feature	Feature coding	Example values	Encoding of example values
Units 1–2	Color of wall on left	Unary coding, 2 units	Black	1 0
			White	0 1
Units 3–14	Length of wall on left	Thermometer coding, 12 units	4	1 1 1 1 0 0 0 0 0 0 0 0
			6	1 1 1 1 1 1 0 0 0 0 0 0
			8	1 1 1 1 1 1 1 1 0 0 0 0
			12	1 1 1 1 1 1 1 1 1 1 1 1
Units 15–18	Angle between left and right wall	Unary coding, 4 units	90°	1 0 0 0
			145°	0 1 0 0
			35°	0 0 1 0
			45°	0 0 0 1
Units 19–22	Feature 1 of object placed at location	Unary coding, 4 units	Value 1	1 0 0 0
			Value 2	0 1 0 0
			Value 3	0 0 1 0
			Value 4	0 0 0 1
Units 23–26	Feature 2 of object placed at location	Unary coding, 4 units	Identical to codes used for 4 different values of Feature 1	
Units 27–30	Feature 3 of object placed at location	Unary coding, 4 units	Identical to codes used for 4 different values of Feature 1	
Units 31–42	Length of wall to right	Thermometer coding, 12 units	Identical to codes used for length of wall to left	
Units 43–44	Color of wall to right	Thermometer coding, 2 units	Identical to codes used for length of wall to left	

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

of individual features – all of the representations in Table 1 are elemental in this sense. A configural representation is accomplished by adding additional units that (locally) represent the presence of unique combinations of feature elements (e.g., a single unit that is turned on when a particular collection of different features is present). The representational scheme detailed in Table 1 is not configural in this sense.

Clearly, there are a number of options that could be used to represent the reorientation task – or any other task – for a perceptron. Features can be represented in local or distributed fashions (with many options within each general type of feature) and can be elemental or configural (Dawson, 2008, Chapter 9). The current representation was chosen because the input units represented all of the required information in an extremely simple format, because we have had prior success with combinations of local and distributed representations (Dawson & Zimmerman, 2003), because we have had prior success with elemental encodings (Dawson, 2004, 2008; Dawson & Spetch, 2005), and because pilot simulations indicated that our encoding choice was satisfactory. Obviously, though, other input representations could be used as is the case for <u>any</u> computer simulation. For example, in addition to exploring local/distributed and elemental/configural distinctions, one could explore the use of other features (Cheung et al., 2008; Sturzl et al., 2008), or even the use of features available at more than one corner location simultaneously.

2.2. Defining the reorientation task for a perceptron: responses

When perceptrons are used to model associative learning experiments, it is typical to use the desired response of the perceptron to indicate whether the perceptron was being reinforced for the presented stimulus; this is also standard practice in studies of reinforcement learning (Dawson, 2008; Gluck & Bower, 1988; Gluck & Myers, 2001; Sutton & Barto, 1981, 1998). This is because reinforcement is equated to the presentation of the unconditioned stimulus, which has the function of activating the desired response. All of the perceptrons studied below had a single output unit that used the logistic equation to convert incoming signals into activity that could range between 0 and 1 (for mathematical details, see Dawson, 2004, Chapter 9). If the perceptron was presented information about a location that was reinforced, then the perceptron was trained to turn "on" to the stimulus (i.e., ideally generate an activation of 1). If the location was not reinforced, then the perceptron was trained to turn "off" to the location's information (i.e., ideally generate an activation of 0).

Output unit activity represents the perceptron's judgment of the likelihood that a particular location is associated with reinforcement (Dawson, Dupuis, Spetch, & Kelly, 2009); the higher the activity, the higher the likelihood. Indeed, Dawson et al. (2009) defined a successful operant learning procedure in which perceptron outputs are used as probabilities that determine whether a perceptron would choose to investigate a particular location (and to only be reinforced – or not – when the location is chosen).

2.3. Training a perceptron on the reorientation task

Once a set of desired stimulus-response pairings have been used to define a problem, a learning rule is used to modify a perceptron's connection weights in order to produce a desired stimulus-response mapping. The learning rules that are typically used to train a perceptron are error-correcting (Rosenblatt, 1962). The magnitude of response error is used to modify connection weights in such a way that the next time the stimulus is presented, the perceptron will generate smaller error. Learning ends when the perceptron generates an acceptably small error to every stimulus in the training set. The gradient descent learning rule (Dawson, 2004) was the particular error-correcting algorithm that was employed in the simulations below. All of the simulations were conducted with the Rosenblatt program (Dawson, 2004, 2005), which is available as freeware from the following website: http://www.bcp.psych.ualberta.ca/~mike/Software/Rosenblatt/index.html.

We now turn to describing some of the results that we have obtained when training perceptrons to solve the reorientation task.

3. Simulation 1: wall color cues in rectangular arenas

Geometric properties are those that are described by relational properties (Gallistel, 1990). For example, the geometric information that describes Location 4 in Fig. 1 is "a 90° corner with the long wall on the left and the short wall on the right". However, when geometric cues are all that are available in rectangular arenas like the one in Fig. 1, target locations are ambiguous because more than one location can possess identical geometric properties. In order to uniquely identify a location in the arena, additional information must be supplied. One technique that has been used to provide such information is to make one of the walls in a rectangular arena a unique color (Cheng, 1986). The purpose of the first simulation was to explore perceptron learning of the reorientation task under analogous conditions.

3.1. Method

Perceptrons were trained on the reorientation task in a 8×4 arena in which the left wall in Fig. 1B was coded as having one color ("white"), while the remaining three walls were coded as having a second color ("black"). No objects were present at any time in this simulation study. Ten different perceptrons each served as an independent "subject" in the experiment. Prior to training, the connection weights were randomized in the range from -0.10 to $+0.10$, with the bias of the output unit's activation function initialized to 0. All training was conducted with the gradient descent method described earlier, with a learning rate of 0.10. The perceptron was reinforced to Location 4, and was not reinforced to the other three locations. Perceptron training was stopped when a "hit" was generated for every training set member. A "hit" was defined as output activity of 0.10 or less for a nonreinforced location, and as output activity of 0.90 or higher for the reinforced location.

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

At the end of training, a perceptron's response to each arena location was recorded. The responses of the perceptron were then examined in three additional arenas (of identical shape and size) without additional training. In the first additional arena the "white" wall was now to the right of the reinforced location (Fig. 1C). In the second, all four arena walls were "black" (Fig. 1A). In the third, all four arena walls were "white".

3.2. Results and discussion

3.2.1. Perceptron responses

Perceptrons converged after 679.70 epochs of training. The first column of data in the upper half of Table 2 presents the mean response of the 10 perceptrons to each of the four arena locations at the end of training. In all cases, the standard deviation of the mean response was less than 0.01. The table labels each of the four arena locations according to reorientation task conventions. Location 4 (e.g., Fig. 1A) was always reinforced during training, so it is called the correct location. From this it follows that Location 1 is the near location, Location 3 is the far location, and Location 2 is the rotational location, because if it is selected this indicates a rotational error.

The first column of data in Table 2 reveals that training was successful: perceptrons turn on to the correct location by producing a mean response of 0.90, and turn off to all nonreinforced locations by producing a mean response of 0.10 or less. Interestingly, the pattern of "off" responses in this column suggests that there may be a combined influence of geometric cues and feature information. First, while perceptrons do turn off to the rotational location, they still generate a stronger response to it (0.10) than to the other two nonreinforced locations. This is because while the wall color provides information to identify this corner as one that is not reinforced, its geometric cues are identical to those of the reinforced corner.

A stronger test of the potential for combined effects of geometry and features can be conducted by placing geometric and local cues in conflict (Brown et al., 2007; Gray, Bloomfield, Ferrey, Spetch, & Sturdy, 2005; Sovrano, Bisazza, & Vallortigara, 2007). This manipulation is represented in the second column of means in Table 2. When the local feature is moved to a different place in the arena, the perceptron generates its strongest response (0.49) to the previously nonreinforced near location, because the "white" wall is now to the left of this location. The response of the perceptron is attenuated, though, by the fact that the geometry of near location is different from that of the previously reinforced correct location. Indeed, the geometry of the rotational location still causes the perceptron to generate activity of 0.10, which is 10 times higher than it generates to the far location, which has neither the correct geometric or feature information. Finally, the perceptron generates a response of 0.08 to the correct location, because while its geometry is correct, the fact that it now has a white wall on its left is a signal that this location will not lead to reinforcement.

Another approach to examining the contribution of geometric cues to perceptron activity is to take the trained perceptrons and observe their responses when they are placed in a rectangular arena whose walls are all "black", which is a color that never reliably identified the reinforced location. It can be seen from the third column of means in Table 2 that when placed in the all "black" arena, the perceptrons turned "off" to all four locations. However, variations in these "off" responses reflect the influence of geometry. In particular, the average response to the two locations with correct geometry (correct and rotational) was 0.10, which was 10 times larger than responses generated to the other two locations that had the wrong geometry (i.e., short wall on the left and long wall on the right).

The post-testing of animals in arenas in which none of the walls have the identifying feature has been reported in the literature. However, one manipulation that has rarely been examined is one in which post-testing occurs in an arena in which all of the walls have the identifying feature (Batty, Bloomfield, Spetch, & Sturdy, 2009). It is extremely easy to explore perceptron responses to this novel arena. The final column of means in Table 2 reports the results of this condition. On the one hand, these results are proportionally similar to those of the all-"black" arena: the strongest responses are to the geometrically equivalent correct and rotational locations, and the weakest responses are to the geometrically incorrect near and far locations. On the other hand, the responses in this arena are markedly and quantitatively different from those of the all-"black" arena: the perceptrons generate responses of 0.88 to the geometrically correct locations, and responses of 0.44 to the other two locations. Again, these results reflect the combined influence of geometric and feature information. In this arena, all locations possess the unique identifier of reinforcement (a "white" wall to the left). When this is combined with correct geometry, perceptron re-

Table 2

Average responses of perceptrons to each arena location in Simulation 1. Location 4 was the reinforced location (i.e., the correct corner).

Response type	Arena location	Arena condition			
		Training (white wall on left)	White wall on top	All walls black	All walls white
Perceptron activity	Near	0.01	0.49	0.01	0.44
	Rotational	0.10	0.10	0.10	0.88
	Far	0.01	0.01	0.01	0.44
	Correct	0.90	0.08	0.10	0.88
Choice rule	Near	0.01	0.72	0.05	0.17
	Rotational	0.10	0.15	0.45	0.33
	Far	0.01	0.01	0.05	0.17
	Correct	0.88	0.12	0.45	0.33

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

sponses become very high. However, they are not quite as high as those observed in the training arena, because neither the correct nor the rotational location has a "black" wall on the right in the all "white" arena. The high responses to the near and far locations, which have incorrect geometry, are being produced mostly by the presence of the "white" wall to the location's left.

Why are perceptron responses so attenuated in the all "black" arena, and so amplified in the all "white" arena? This result reflects the amount of information provided by the various cues during training. In the training arena, having a "white" wall to the left of a location was a perfect signal for reinforcement. Black walls to the right signaled reinforcement for some locations, but not others, while black walls on the left always signaled no reinforcement. Thus in the all "black" arena, the local features for the most part signaled the absence of reinforcement, which dramatically attenuated perceptron responses. In contrast, in the all "white" arena, the local cues at all corner ("white" on the left) were strong cues for reinforcement, causing much larger responses in the perceptrons.

This interpretation is borne out by examining the connection weights for the two units that represent the color of the wall on the left and the two units that represent the color of the wall on the right. For each of these pairs of input units, the first unit codes "black", and the second unit codes "white". At the end of training a randomly selected "subject" network had connection weights of −2.21 and 2.20 for the two left wall units, and had connection weights of 0.01 and −0.19 for the two right wall units. This pattern of connections indicates two things. First, the color of the wall on the left was far more important than the color of the wall on the right, because the absolute values of the first pair of weights are much greater than those of the second pair. Second, "black" on the left was an inhibitory cue (indicated by its strong negative weight) while "white" on the left was an excitatory cue (indicated by its strong positive weight). The fact that the networks have learned both inhibitory and excitatory cues, as revealed by these connection weights, explains the different levels of perceptron responses when all arena walls are one color versus the other.

3.2.2. *Converting perceptron responses to choice probabilities*

The results of Simulation 1 were discussed above by reporting perceptron responses to the cues available at each of the four locations. Each of these responses can be interpreted as the perceptron's prediction of the probability of receiving a reward at each location, where this probability is computed independently for each location (Dawson et al., 2009). However, in the reorientation task literature, it is typical to report data in terms of the probability of choosing one location over the other three. In order to strengthen the connection between our simulations and this literature, can perceptron responses be converted into choice probabilities?

There are many different ways in which this could be done. We elected to employ an equation which has a long history in mathematical psychology (Gulliksen, 1953; Herrnstein, 1970; Luce, 1959, 1961, 1977; Thurstone, 1930). For instance, Thurstone (p. 470) defined the probability of an act leading to successful consequences (P) as $P = s/(s + e)$, where s is the strength of a successful response and e is the strength of competing responses. Following the form of this equation, perceptron responses can be converted into choice probabilities by dividing each (average) response by the sum of the (average) perceptron responses to all of the arena locations. The lower half of Table 2 reports the results of Simulation 1 after the perceptron responses in the upper half of the table were converted into choice probabilities using this simple rule.

For the most part, the choice probabilities reported in the lower half of Table 2 mirror the preceding discussion of perceptron responses. After training, there is an overwhelmingly large probability of choosing the correct location (first column of data). When the wall color feature is moved to a different location, choice probability shifts with the feature (second column of data).

One interesting result is evident when choice probabilities are computed for the condition in which the feature is removed from the arena. On the one hand, strong rotational error is evident: there is a 0.45 probability of choosing the correct location, and a 0.45 probability of choosing the rotational location (third column of data). This type of result is what one would expect from surveying the literature (Cheng & Newcombe, 2005). On the other hand, the choice probabilities computed for this condition hide the fact that perceptron responses are very weak for every location in this condition. This suggests that a complete account of reorienting behavior may require reporting more than one dependent measure, both in simulations and in experimental studies of biological agents. Choice probabilities might reflect the relative attractiveness of the various locations, but other measures – such as the latency to choose a location to explore – might be an indication of the absolute attractiveness of a location, which in our simulations is revealed by perceptron responses.

Another interesting result is evident when the responses to the "all white" arena are converted into choice probabilities (fourth column of data). In this case, even though there is an increase in perceptron responses to the correct and rotational locations (relative to the "all black" arena), the choice probabilities for these two locations actually decrease to 0.33. This is because the increase in perceptron responses to the near and far locations increases the choice probability to 0.17 for each. Again, this result suggests that reorientation task behavior may require measures of both absolute and relative attractiveness of locations.

3.3. *Summary*

One of the general findings in studies that have employed features that, by themselves, are sufficiently informative to uniquely identify a target location, is that animals do not purely rely on these features. Instead, geometric cues are also attended to, even though these cues are unnecessary. The results from Simulation 1 indicate that a similar situation is true for perceptrons, when either perceptron responses or choice probabilities are considered. In order to explain the patterns of responses of perceptrons to the different locations, one must appeal to a combination of geometric and feature cues. Furthermore, these results suggest that if one only examines choice probabilities, then some of the mechanisms that mediate task performance might not be revealed. This is because there were some intriguing differences when results were reported in terms of perceptron responses and when these responses were converted into choice probabilities.

4. Simulation 2: reorientation in kite-shaped arenas

Reorientation has been studied in arenas of various shapes and sizes. In Simulation 2, we consider kite-shaped arenas that can been used (Graham, Good, McGregor, & Pearce, 2006; Pearce, Good, Jones, & McGregor, 2004; Pearce, Graham, Good, Jones, & McGregor, 2006), in combination with manipulation of wall color features, to explore some of the predictions that emerge from the theory that arena shape is processed by a geometric module (Cheng, 1986; Gallistel, 1990). Fig. 3) illustrates some example kite-shaped arenas that were explored in this simulation. These arenas were similar to the rectangular arenas that were explored in Simulation 1 in the sense that the long sides of the kite were twice the length of the short sides. The difference between the two types of arenas is defined by the angles at each vertex of the kite: while the angle at locations 2 and 4 is still 90°, the angle at location 1 is 145° and the angle at location 3 is 35°. This is important, because the combination of side lengths and vertex angles ensures that geometric features uniquely define each of the locations, indicating that geometry can be used to accomplish the reorientation task without relying upon additional cues. For instance, Location 4 is uniquely defined as the location with a long wall on the left meeting a short wall on the right at an angle of 90°, while Location 2 is uniquely defined as the location with a short wall on the left meeting a long wall on the right at an angle of 90°.

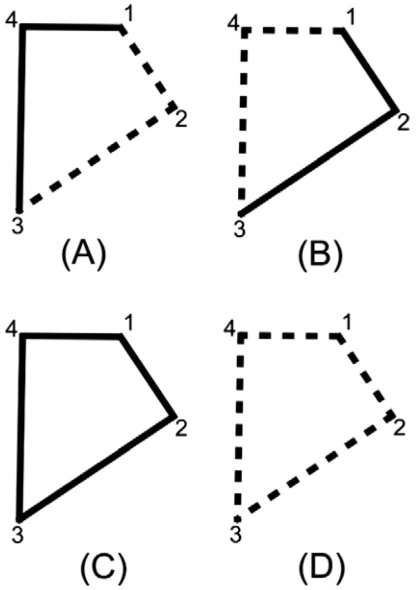

Fig. 3. Four different types of kite-shaped arenas used in Simulation 2. Solid lines indicate black walls, and dashed lines indicate white walls. The numbers correspond to the different locations for which features were presented to the perceptrons.

One study trained rats to navigate to goal locations in a kite-shaped Morris water maze (Graham et al., 2006). Three different conditions were examined. In the shape + -color condition, the rats were reinforced at Location 4, which in training was always bounded by black walls and the opposite two walls were white (Fig. 3A). Thus the goal location was defined by both unique geometry and by unique local features. In the shape condition, the rats were always reinforced at Location 4, but in some trials it was bounded by black walls (Fig. 3A) while in others it was bounded by white walls (Fig. 3B). Thus, in this condition the goal location was defined by geometry and not by local features. In the color condition, the rats were trained in the same two types of arenas that were used in the shape condition. However, they were reinforced at Location 4 when it was defined by black walls, and were reinforced at Location 2 when it was defined by black walls. Thus, in this condition the goal location was only defined by local features (two adjacent black walls), and was not defined by geometric shape.

Graham et al. (2006) divided the kite-shaped arena into two large and two small quadrants, where the quadrant containing Location 2 and the quadrant containing Location 4 each contained 38.2% of the arena area. After training, animals in each of the three conditions were tested in a kite-shaped arena in which all of the walls were black. Animals trained in the shape + color condition spent more time in the quadrant that contained the target location than did animals trained in the other two conditions. The animals were also tested in a square arena in which two adjacent walls were black, and the other two adjacent walls were white. Note that in this arena the unique geometric cues that define each location in the kite-shape no longer exist. It was found that animals trained in the shape + color condition and in the color condition had a strong preference to stay in the quadrant that contained the two black walls. These results indicated that the rats in the shape + color condition learned about both shape and color, even though a geometric module could have solved the task without requiring any information about wall color.

The purpose of the second set of simulations was to explore the behavior of perceptrons when trained under conditions that simulated the three different training conditions studied by Graham et al. (2006). Importantly, the purpose of Simulation 2 was not to attempt to model the results of Graham et al. (2006). One reason for this is that it is difficult to map perceptron responses onto the measures that Graham et al. employed. For instance, Graham et al. measured the percentage of time that animals spent in fairly large quadrants of the water mazes, and did not report measures related to the smaller quadrants that contained Locations 1 and 3 in arenas like those illustrated in Fig. 3.

4.1. Method

Perceptrons were trained in kite-shaped arenas whose long sides were eight arbitrary units long and whose short sides were four arbitrary units long. Each arena had two walls of one color (e.g., black) flanking the correct location

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

(Location 4 in Fig. 3A), and two walls of the opposite color (e.g., white) flanking the rotational location (Location 2 in Fig. 3A). Wall lengths, angle between walls, and wall color were coded using the encoding scheme described earlier in Table 2. As was the case in Simulation 1, no objects were present at any time in this simulation study, so all of the object property inputs were turned off.

In order to simulate the different training conditions, each epoch of training involved training a perceptron in two different arenas (i.e., the training set consisted of eight different training patterns, four for one arena, and the other four for a second arena). For the shape + color, both of the arenas had black walls flanking the correct location (Fig. 3A), and the perceptron was reinforced for this location (and not reinforced at the other three locations) in both arenas. For the shape condition, one of the arenas had black walls flanking the correct location (Fig. 3A), while the other had white walls flanking the correct location (Fig. 3B). For both of these arenas, the perceptron was only reinforced at this location. For the color condition, the perceptron was reinforced at the correct location when it trained on the arena that had black walls flanking this location. The perceptron was reinforced at the rotational location in the other arena, which had black walls flanking Location 2 and white walls flanking Location 4. Ten different perceptrons served as subjects in each of these three conditions.

After perceptrons were trained to convergence, their responses were recorded, without additional training, in four different kite-shaped arenas: all black walls (Fig. 3C), all white walls (Fig. 3D), black walls surrounding Location 4 and white walls surrounding Location 2 (Fig. 3A), and white walls surrounding Location 4 and black walls surrounding Location 2 (Fig. 3B).

4.2. Results and discussion

On average, the perceptrons in the shape + color condition converged (generated a “hit” to each of the eight locations in the training set) after an average of 145.10 epochs of training; in the shape condition an average of 190.40 epochs of training was required; in the color condition an average of 281.00 epochs of training was required.

4.2.1. Perceptron responses

The top part of Table 3 presents the average response of perceptrons trained in the different conditions to four different test kite-shaped arenas. The “Black Kite” column of Table 3 is one of the test conditions examined by Graham et al. (2006). One of the conclusions that they drew from their results is that animals trained in the shape + color condition were affected by color, even though shape was sufficient to accomplish the navigational task. Table 3 supports a similar interpretation. Note that when all of the kite walls are black, the perceptron’s responses to the near, rotational, and far locations are noticeably higher than the responses in the condition where black walls flank the correct location (and white walls flank the rotational location, the third column of means in Table 3). This increase in response can only be attributed to a difference in wall color, because arena geometry is constant in all of the test conditions. Thus shape + color perceptrons learned about both of these features. Similarly, note that when tested in the black kite, perceptrons trained in the color condition generate weaker responses to the near and far locations than to the correct and rotational locations. This suggests that they learned information about geometry in addition to color. This is because while they were trained

Table 3
Average responses of perceptrons to kite-shaped arenas after training in Simulation 2.

Response type	Training condition	Arena location	Arena condition			
			Black kite	White kite	Black and white kite (black flanking correct)	Black and White kite (white flanking correct)
Perceptron activity	Shape + color	Near	0.23	0.03	0.08	0.10
		Rotational	0.10	0.01	0.01	0.10
		Far	0.17	0.02	0.07	0.06
		Correct	0.90	0.49	0.90	0.49
	Shape	Near	0.08	0.08	0.08	0.08
		Rotational	0.02	0.02	0.02	0.02
		Far	0.07	0.07	0.07	0.07
		Correct	0.90	0.90	0.90	0.90
	Color	Near	0.49	0.01	0.07	0.07
		Rotational	0.90	0.05	0.05	0.90
		Far	0.47	0.00	0.06	0.06
		Correct	0.90	0.05	0.90	0.05
Choice rule	Shape + color	Near	0.16	0.05	0.08	0.13
		Rotational	0.07	0.02	0.01	0.13
		Far	0.12	0.04	0.07	0.08
		Correct	0.64	0.89	0.85	0.65
	Shape	Near	0.07	0.07	0.07	0.07
		Rotational	0.02	0.02	0.02	0.02
		Far	0.07	0.07	0.07	0.07
		Correct	0.84	0.84	0.84	0.84
	Color	Near	0.18	0.09	0.06	0.06
		Rotational	0.33	0.45	0.05	0.83
		Far	0.17	0.00	0.06	0.06
		Correct	0.33	0.45	0.83	0.05

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

to respond in two geometrically different locations (when flanked by the target color), these locations always shared a geometric property (a 90° angle between walls) that is not shared by these other two locations.

While the results of the black kite test condition supports an important conclusion drawn by Graham et al. (2006), there are some striking differences between the perceptron data and the animal data. Graham et al. noted that the animals in the shape + color condition spent significantly more time in the arena quadrant that contained the correct location than they did in the quadrant that contained the rotational location; animals trained in the other two conditions tended to spend similar amounts of time in both quadrants. Again, we point out that it is difficult to map perceptron response strength to a particular location to this dependent measure of Graham et al. Nevertheless, the black kite results column in Table 3 would seem to indicate that there should be little difference between the shape + color and the shape conditions, given that both indicate a very strong response to the correct location. (That the color condition perceptrons give equally strong responses to the correct and rotational locations is arguably more consistent with the Graham et al. results.)

The fact that both shape + color and color perceptrons have an equal preference for the correct location, while different from the observations of Graham et al. (2006), does not rule out a perceptron-like account of their data. In Simulation 2, all perceptrons were trained to equally high performance in the reorientation task. However, in the first experiment of Graham et al. the animals trained in the shape condition never learned to perform the task as well as did animals in the other two conditions. It could be that differences between animals in their two conditions might reflect different degrees of learning, and that if learning was equated, their results might be more similar to those reported in the black kite column of Table 3.

Perhaps of more interest is the performance of perceptrons in test arenas that were not studied by Graham et al. (2006). For example, in the shape condition, color was never a reliable predictor of reinforcement. So it is not surprising that the responses of this perceptron to all kite locations are identical, regardless of the colors of the walls (white, black flanking the correct location, black flanking the rotational location). Similarly, in the color condition, color was the only reliable predictor, so again it is not surprising that perceptron responses are highest when the color cues are appropriate, and are near zero when these cues are not appropriate. However, if the differences between shape + color and shape conditions are critical for testing the geometric modularity assumption, then it would appear that the strongest differences between these conditions are realized when the color cues are removed from the arena (white kite), or when the color cues are in conflict with geometric cues (black flanking the rotational location). In these two conditions, the geometry surrounding the correct location is correct, but the color cues that the shape + color condition experienced are not. As a result, the shape + color perceptrons generate a response that is nearly half of the response that they generate when color cues are correct, and nearly half of the response of the shape perceptrons to the same locations.

Even though the shape + color perceptron responses to the correct location are attenuated when the black walls are removed from the arena, or when they are moved to a location that is in conflict with learned geometric cues, it is still the case that the correct location generates a much stronger response than do any of the other three locations. It is not clear how this weaker – but still strong – response would be reflected in the dependent measures that were used by Graham et al. (2006). One issue facing associative models of learning in general (Rescorla & Wagner, 1972), and perceptrons as an example of such theories, is the translation of model variables (e.g., associative strengths, connection weights, processing unit activities) into measurable behavior (Dawson, 2008, Chapter 8). The results for perceptrons responding to the kite-shaped arenas that were not tested by Graham et al. indicate that these arenas might be particularly useful for evaluating associative models of the reorientation task under an appropriate mapping of output unit activities to animal behavior.

4.2.2. Choice probabilities

The bottom part of Table 3 reports the results of these simulations after perceptron activities were converted to choice probabilities using the same procedure that was described in Simulation 1. In general, an examination of the choice probabilities supports the interpretation that was given earlier when perceptron activities were discussed. However, there are some interesting quantitative differences between the two measures. Consider the shape + color condition when perceptrons were tested in the white kite. In terms of perceptron activity, the correct location generated a response of 0.49. However, when this is converted to a choice probability, it increases to 0.89 because of the low activities associated with each of the other three locations. This again indicates that choice probability does not necessarily provide an accurate estimate of responses to cues present at locations considered independently, and suggests that empirical studies of reorientation might be improved by reporting multiple dependent measures (e.g., choice probability as well as latency to respond, or the actual routes taken by an agent as it reorients to the new arena (Wystrach & Beugnon, 2009)).

5. Simulation 3: distinct objects in corners of rectangular arenas

Another approach to providing feature information in the reorientation task is to provide distinctive objects to distinguish the target location from others (Kelly et al., 1998). This third simulation adopted this paradigm. During training, a three-featured "object" was placed at each of the four arena locations in a rectangular arena. After training, the behavior of the perceptron was observed in additional arenas in which the objects were moved to different locations, or were removed.

5.1. Method

Perceptrons were trained in a rectangular arena with a length of eight arbitrary units, and a width of four arbitrary units. The perceptron was reinforced when presented the stimulus information from Location 4 in Fig. 1 (the correct location) and was not reinforced for the other three locations. A unique object, each defined by three features, was present at each corner of this arena during training (Fig. 1D). An object's features were arbitrary; they were vectors of numbers that were thought of as representing some visual properties that could be detected by an agent (e.g., shape, color, height). Each feature could have one of four different values. In Simulation 3, each object was defined by three feature values that were unique – that is, the features possessed by one object were not possessed by the other objects in the arena. The feature values for the object at the near location (Location 1) were $f_1 = 1$, $f_2 = 1$, and $f_3 = 1$. Similarly, the values for f_1, f_2 and f_3 were all 2 for the object at the rotational location (Location 2), were all 3 for the object at the far location (Location 3), and were all 4 for the object at the correct location (Location 4). Each feature value was encoded as indicated earlier in Table 1. Additional input units were used to represent wall length, wall color, and the angle between walls using the same encoding scheme that was described in Table 1, and employed in Simulations 1 and 2.

Ten different perceptrons were trained to convergence using the same procedure that was described for Simulation 1. After training, the responses of perceptrons to four different arena configurations were examined. The first was the arena in which the perceptron was trained. The second was the same arena, with each object moved to a new location, the nearest corner in a clockwise direction. That is, the object at Location 1 in Fig. 1D was moved to Location 2, the object at Location 2 was moved to Location 3, and so on, producing the arrangement of objects in Fig. 1E. This change in object locations is formally equivalent to an affine transformation; first the arena is compressed so that its two long sides are half of their original length; second the two short sides (prior to the first transformation) are stretched to twice their original length (Gallistel, 1990, p. 186). Such an affine transformation places geometric and feature cues in conflict because after the transformation the previously reinforced location is presented with the correct geometry, but with an incorrect object and a previously nonreinforced location (Location 1) is presented with a correct object, but incorrect geometry. The third arena that was examined was identical to the training arena, except that the objects at the correct and rotational locations were removed (by setting the activity of all 12 object-feature units to 0). This arena was studied because it has been shown that pigeons (Kelly et al., 1998) can use objects at nonreinforced locations as landmarks to find the reinforced locations, although most species (rats, chicks, fish) do not (Cheng, 1986; Sovrano et al., 2003; Vallortigara, Zanforlin, & Pasti, 1990). The final arena that was studied was one in which all four objects were removed. This arena is of interest because it can be used to assess a perceptron's sensitivity to geometric cues when no feature cues are available.

5.2. Results and discussion

On average, the perceptrons converged after 252.2 epochs of training.

5.2.1. Perceptron responses

The mean responses of the perceptrons after training to each of the four locations in four different arena settings are provided in the upper part of Table 4. The standard deviations of each of these means are 0.02 or smaller.

The first column of data in Table 4 reveals that training was once again successful: perceptrons turn on to the correct location by producing a mean response of 0.90, and turn off to all nonreinforced locations by producing a mean response of 0.09 or less. As was the case in Simulation 1, the pattern of "off" responses in this column suggests that there may be a combined influence of geometric cues and feature information, because the perceptrons generate a stronger response to the rotational location (0.09), which has correct geometry, than to the other two nonreinforced locations.

The influence of featural and geometric clues on perceptron responses is also supported by examining the connection weights in trained networks. In a typical network, connection weights for any of the three units representing a value of 4 for one of the three features have excitatory weights whose values are usually in the range of 0.80 (because this feature value is a reliable cue for reinforcement at the correct location); the three units representing a value of 2 for any of the three features have inhibitory weights whose values are usually in the range of −0.65 (because this feature value is a reliable cue for no reinforcement at the rotational location); the remaining units that represent other feature values have markedly smaller

Table 4
Average responses of perceptrons to each arena location in Simulation 3. Location 4 was the reinforced location (i.e., the correct corner).

Response type	Arena location	Arena condition			
		Objects in original locations	Objects moved to location on right	Objects removed from rotational and correct locations	All objects removed
Perceptron activity	Near	0.04	0.45	0.04	0.06
	Rotational	0.09	0.29	0.40	0.40
	Far	0.03	0.01	0.03	0.06
	Correct	0.90	0.29	0.40	0.40
Choice rule	Near	0.04	0.43	0.05	0.07
	Rotational	0.08	0.28	0.46	0.43
	Far	0.03	0.01	0.03	0.07
	Correct	0.85	0.28	0.46	0.43

weights. However, in the same network, weights that are involved in coding a wall length of 8 (i.e., units 7 through 10 for the left wall and units 35 through 38 for the right wall) also have developed moderately strong weights whose absolute value is typically from 0.3 to 0.4. Interestingly, these length units for the left wall have excitatory connection weights, and the same units for the right wall have inhibitory weights, indicating that the network is sensitive to the fact that at the correct location the left wall is longer than the right wall.

The second column of means in Table 4 examines the response of the perceptrons to an arena which has undergone an affine transformation, and in which geometric and feature cues are in conflict. Perceptron responses indicate that both geometric and feature cues were processed, even though geometric cues were not necessary to learn the location of reinforcement. First, when the correct feature cues are moved to the near location, the perceptron's response to this location increases from 0.04 in the original arena to 0.45. However, a strong response is still evident for the correct location (0.29), even though incorrect features are present at this location. This strong response must be mediated by the correct geometry at this location. This is supported by the rotational error of the perceptron, in which an equally strong response is generated to the rotational location (0.29). This response is the result of correct geometry combined with a change in feature information that was inhibiting responses to this location after initial training. The weakest response is to the far location, which has both incorrect featural and incorrect geometric information.

Rotational error is also evident in the final two conditions, represented by the final two columns of means in Table 4. When feature information is removed from the correct and rotational locations, moderate responses are generated by the perceptron to both of these two locations. This illustrates rotational error. The fact that the column 3 means are almost identical to those of column 4 reflects the fact that perceptrons are unable to use objects at non-reinforced locations as landmarks for locating the reinforced corner. The fact that rotational error is evident in both of these columns again shows that the perceptrons did process geometric cues, even though the feature cues by themselves were capable of identifying the reinforced location during learning. This result is important, because it suggests that the feature cues did not completely overshadow learning about the geometric cues in this simulation. Several studies have shown that feature cues near a reinforced location usually do not prevent the learning of geometric cues (Hayward, Good, & Pearce, 2004; Pearce, Ward-Robinson, Good, Fussell, & Aydin, 2001; Wall, Botly, Black, & Shettleworth, 2004). When overshadowing occurs in a perceptron (e.g., Dawson, 2008, Chapter 4), there is usually no response at all to the overshadowed cues, which certainly is not the case in Table 4.

5.2.2. Choice probabilities

The lower part of Table 4 reports the results of Simulation 3 after perceptron activity has been converted into choice probability using the same procedure that was used in the previous simulations. The choice probabilities support the same interpretation that was discussed in the analysis of perceptron activity, and while there are some quantitative differences between the upper and lower parts of Table 4, these differences are fairly minor.

6. Simulation 4: unique, but less salient, feature cues

When objects (often colored and patterned panels) are used to provide feature cues in the reorientation task, the goal is usually to provide a unique visual marker at each arena location. With a perceptron we can easily manipulate the salience of featural cues during training, as is shown in Simulation 4. It is identical to Simulation 3, with the exception that the objects at each location have only one distinctive feature, and share the other two features with every other landmark in the arena.

6.1. Method

The method was identical to Simulation 3, with the exception of the features that were used to define the four different objects at each of the Fig. 1 locations. In Simulation 4, the values for f_1 were unique at each location (equal to 1 at Location 1, 2 at Location 2, and so on), while the values for f_2 and f_3 were equal to 2 for all four objects. Thus, each object has one unique feature (f_1), but shares the same values for the other two features.

6.2. Results and discussion

On average, the perceptrons in this simulation converged after 717.8 epochs of training. In other words, by having less-salient objects, learning in Simulation 4 took over 1.4 times the learning than was required in Simulation 3.

6.2.1. Perceptron responses

The mean responses of the perceptrons after training to each of the four locations in four different arena settings are provided in the upper part of Table 5. The standard deviations of each of these means are 0.02 or smaller.

A comparison of Table 5 to Table 4 indicates that the qualitative pattern of results for Simulation 3 was identical to that observed in Simulation 2, indicating that once again both geometric and feature cues were being used to solve the reorientation problem. However, there are important quantitative differences between the two tables. First, when perceptrons are placed in the arena that has undergone an affine transformation (the second column of means in Table 5) their average response to the near location is 0.24, which is nearly half the response that was observed to the same location in Simulation 3. Second, for this same arena, the responses to the rotational and correct locations are higher in Simulation 4 than was the case in Simulation 3. After the affine transformation, both of these locations are associated with correct geometric cues and incorrect feature cues. Third, the perceptron responses to these same two locations when two objects are removed, or when all four objects are removed, are again higher in Simulation 4 than in Simulation 3. In these cases, geomet-

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

Table 5
Average responses of perceptrons to each arena location in Simulation 4. Location 4 was the reinforced location (i.e., the correct corner).

Response type	Arena location	Arena condition			
		Objects in original locations	Objects moved to location on right	Objects removed from rotational and correct locations	All objects removed
Perceptron activity	Near	0.02	0.24	0.02	0.03
	Rotational	0.10	0.39	0.49	0.49
	Far	0.02	0.00	0.02	0.03
	Correct	0.90	0.40	0.49	0.49
Choice rule	Near	0.02	0.23	0.02	0.03
	Rotational	0.10	0.38	0.48	0.47
	Far	0.02	0.00	0.02	0.03
	Correct	0.87	0.39	0.48	0.47

ric cues are correct, and no feature cues are present. Combined, the quantitative differences between Tables 5 and 4 indicate quite clearly that geometric cues were more powerful than feature cues in Simulation 4, which is sensible given that the intent of Simulation 4 was to provide feature cues that, though still unique, were less salient than those studied in Simulation 3.

A comparison of the connection weights for a typical Simulation 4 network with those for a Simulation 3 network supports the above interpretation. First, the general pattern of wall length unit connection weights for the Simulation 4 network was the same as was previously described for Simulation 3. However, the absolute values of these weights in Simulation 4 were often higher by values of 0.1–0.2, suggesting a greater emphasis on geometry when featural information was less salient. In terms of units representing feature values, the unit representing a value of 4 for feature 1 (signaling reinforcement) had a very high excitatory weight (2.0 or higher), while the unit representing a value of 2 for the same feature (signaling no reinforcement at the rotational location) had an extreme inhibitory weight (-2.0 or lower). However, all of the units for features 2 and 3 had near-zero weights. In short, the network placed greater emphasis on the two cues required to differentiate the correct and rotational locations in this simulation, but could not use any other featural cues. This is consistent with the fact that featural cues were less salient in this simulation.

6.2.2. Choice probabilities

The lower part of Table 5 reports the results of Simulation 4 after perceptron activity has been converted into choice probability using the same procedure that was used in the previous simulations. The choice probabilities support the same interpretation that was discussed in the analysis of perceptron activity, and while there are some quantitative differences between the upper and lower parts of Table 5, these differences are fairly minor.

7. Simulation 5: moving landmark features

In Simulations 3 and 4, each of the four arena locations in Fig. 1 were associated with objects that were defined by at least one unique feature, and these cues were put in conflict with geometric cues by moving objects in their entirety. One advantage of our modeling approach is that alternative transformations of object features can be performed. For instance, imagine that agents are not processing landmarks as whole objects, but are rather learning about the information carried by local features of these objects. How would an agent respond if some of the features were moved in one direction, and other features moved in the opposite direction? To our knowledge, this situation has not been studied. However, it is easily simulated with a perceptron, and produces an easily explained – but counterintuitive – result that raises an important issue that various theories of spatial reorientation would have to address if the same pattern of data were observed in biological agents.

In Simulation 5, perceptrons were trained on the reorientation task using landmarks at each location, where each landmark was defined by three unique features. The training phase of this simulation was identical to the training phase in Simulation 3. In particular, the feature values for each object were identical to those described earlier in Simulation 3, as can be seen in Table 6 below. The test phase of this simulation was created by moving the f_1 feature value in a clockwise direction to the nearest location in the test arena, by moving the f_3 feature value in a counterclockwise direction to the nearest location in the test arena, and by keeping the f_2 feature value in the original location. The features that defined these objects are also given below in Table 6. Each of these new landmarks was completely new; that is, none of these feature triplets defined a landmark that was used at any location during training. The question of interest was how these new landmarks would affect perceptron responses in the test arena.

7.1. Method

The training method and training patterns for Simulation 5 were identical to those used in Simulation 3, apart from the transformation of features in the test arena that was described above.

7.2. Results and discussion

On average, the perceptrons converged after 251.9 epochs of training.

7.2.1. Perceptron responses

The mean responses of the perceptrons after training and in the test arenas are provided in upper third of Table

Table 6
Average responses of perceptrons to each arena location in training and test arenas in Simulation 5. Location 4 was the reinforced location.

		Arena			
		Training		Test	
	Arena location	Landmark features f_1, f_2, f_3	Perceptron responses	Landmark features f_1, f_2, f_3	Perceptron responses
Perceptron activity	Near	1,1,1	0.04	4,1,2	0.07
	Rotational	2,2,2	0.09	1,2,3	0.19
	Far	3,3,3	0.04	2,3,4	0.06
	Correct	4,4,4	0.90	3,4,1	0.52
Choice rule	Near	1,1,1	0.04	4,1,2	0.08
	Rotational	2,2,2	0.08	1,2,3	0.23
	Far	3,3,3	0.04	2,3,4	0.07
	Correct	4,4,4	0.84	3,4,1	0.62
Miller/Shettleworth model	Arena location	Landmark features	Choice probabilities	Landmark features	Choice probabilities
	Near	f_1, f_2, f_3	0.01	f_2, f_6, f_{10}	0.20
	Rotational	f_4, f_5, f_6	0.00	f_1, f_5, f_9	0.14
	Far	f_7, f_8, f_9	0.01	f_4, f_8, f_{12}	0.20
	Correct	f_{10}, f_{11}, f_{12}	0.98	f_7, f_{11}, f_3	0.46

6. The standard deviations of each of these means are 0.02 or smaller.

The one intuitive result that is presented in Table 6 is the attenuation of the perceptron's response to the correct location – in the test arena, the response to this location is nearly half of the response to the same location in the training arena. However, the values of f_1 and f_3 at the correct location in the test arena are no longer equal to 4. As these new values were not previously associated with reinforcement, it is not surprising that they cause a reduction in perceptron responses when they are shifted to Location 4.

Two other less intuitive results are also evident in Table 6. First, there is a very small increase in perceptron responses to the near and far locations. During training, none of the three feature values at either of these locations was reinforced. In the test arena, for both locations, one of the new feature values was reinforced. However, its presence does not seem to increase perceptron responses as much as might be expected. Second, during training none of the rotational location feature values had been reinforced. In the test arena, it is again true that none of the rotational location feature values had been reinforced. Nevertheless, perceptron responses to this location in the test arena are nearly doubled relative to responses to the same location in the training arena!

Why is it that, in the test arena, responses to the rotational location are surprisingly high, and responses to the near and far locations are surprisingly low? The answer to this question comes from realizing that during training, not only must the perceptron learn associations that signal that Location 4 is the target, but it must also learn associations that signal that the geometrically equivalent Location 2 is not the target. This is accomplished, during training, by assigning high excitatory associations to the units that represent the feature values that are present at the correct location, and by also assigning high inhibitory associations to the units that represent the feature values that are present at the rotational location. In the test arena, two of these inhibitory feature values have been shifted away from the rotational location. This release from inhibition produces the surprising increase in rotational location responses. As well, when these inhibitory feature values are shifted to the near and far locations in the test arena, they tend to attenuate the excitatory signal that is added to these locations when the feature value "4" is shifted to them as well. As a result, responses to these locations do not increase as much as might be expected.

An interpretation of the connection weights of trained networks can be easily used to support the argument made in the previous paragraph. In a typical network, the three input units that represent the presence of value 4 for any of the three landmark features have connection weights that are strongly excitatory, having a value of 0.80 or higher. Of course, this feature value was associated with reinforcement during training. In contrast, the three input units representing the presence of value 2 for any of the three landmark features have connection weights that are strongly inhibitory, having a value of −0.54 or more negative. This feature value was associated during training with the nonreinforced location that was geometrically equivalent to Location 4. All the remaining feature values are represented by connection weights that are more moderately inhibitory, ranging in value from −0.10 to −0.25. In short, the reorientation task is being solved by making excitatory associations with the features that accompany the correct location, and at the same time by making inhibitory associations with the features that accompany the rotational location.

7.2.2. *Choice probabilities*

The middle third of Table 6 reports the results of Simulation 5 after perceptron activities have been converted to choice probabilities using the same procedure employed in the previous simulations. When this is done, the results follow the same pattern that was revealed when perceptron responses were examined. In particular, training when the features were located in their original positions results in a high preference to choose the correct location, and much lower preference for all other locations (although

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

the preference for the rotational location is slightly higher than for the near or far locations). When features are then moved to their new locations, there is a marked decrease in preference for the correct location, a slight increase in preference for the near and far locations, and a much stronger increase in preference for the rotational location. This latter effect is, in fact, larger for the choice probability data than it is for perceptron response data. Furthermore, this effect must be due to a release from inhibition, because none of the features present at the rotational location in the test phase had been previously reinforced.

7.2.3. Miller/Shettleworth predictions

Given that Simulation 5 reports the results of a study that has not yet been conducted with biological subjects, it is premature to consider what it implies for the reorientation task. However, these results do indicate that inhibition plays an important role in how a particular model, the perceptron, develops associations that can be used to perform spatial reorientation. This is crucial, because in other theories of reorientation, inhibition does not play as important a role. For instance, in order to correct a flaw in an associative model of the reorientation task, mathematical assumptions are required that reduce the influence of inhibition (Miller & Shettleworth, 2008). The same is the case in theories that would use similarity metrics – but not dissimilarity metrics – to solve the problem by matching new locations to representations of the appearance of earlier ones (Cheung et al., 2008; Stuerzl et al., 2008). Clearly reorientation experiments that explicitly search for effects like release from inhibition would play an important role in differentiating these different models.

In the absence of experimental data, it is important to note that the release from inhibition effect discovered in the perceptron simulations can be used to distance this type of model from other associative models (Miller, 2009; Miller & Shettleworth, 2007, 2008). To demonstrate this, we modified an example reorientation task that was studied in detail by Miller and Shettleworth (2007). In the original example, feature B was contextual information present at all four arena locations, feature F was present only at the correct location, feature G was correct geometric information present at the correct and rotational locations, and feature W was incorrect geometric information present at the near and far locations. We took this example, but expanded feature F to be 12 different feature cues. Features f_1, f_2, and f_3 were present at the near location, features f_4, f_5, and f_6 were present at the rotational location, features f_7, f_8, and f_9 were present at the far location, and features f_{10}, f_{11}, and f_{12} were present at the correct location. We trained the corrected version of the Miller/Shettleworth model (Miller & Shettleworth, 2008) on this elaborated feature set by coding the presence of a feature with 1 and the absence of a feature with 0. The choice probabilities generated by this model after 26 iterations are provided in the bottom third of Table 6. It can be seen that the Miller/Shettleworth model has an overwhelmingly large preference for the correct location, and near zero preference for all other locations, after training, which is to be expected. However, when features are moved (see Table 6) and choice probabilities are recomputed from the associative strengths produced by the model, a pattern of results that is very different from the perceptron is produced. In particular, in the test situation the rotational location is associated with the lowest choice probability of any of the locations. Presumably this is because the Miller/Shettleworth model has a growing likelihood of visiting the correct location, indicating that its underlying mechanisms will involve increasing the strength of excitatory cues, and are less likely to involve making cues signaling the absence of reinforcement more inhibitory. Indeed, an examination of the graphs of cue strength over time that are reported by Miller and Shettleworth (2007) shows that excitation increases to much higher levels than does the (absolute) value of inhibition. This makes release from inhibition less potent, and produces different results for their model than is observed in the perceptron. These differing predictions suggest an exciting avenue for future empirical work.

8. Simulation 6: effects of arena size with feature cues present

Many researchers have explored the effect of changing arena size on the reorientation task (Chiandetti, Regolin, Sovrano, & Vallortigara, 2007; Hermer & Spelke, 1994; Learmonth, Nadel, & Newcombe, 2002; Learmonth, Newcombe, & Huttenlocher, 2001; Learmonth, Newcombe, Sheridan, & Jones, 2008; Ratliff & Newcombe, 2008; Sovrano, Bisazza, & Vallortigara, 2002; Sovrano et al., 2003, 2007; Sovrano & Vallortigara, 2006; Vallortigara, Feruglio, & Sovrano, 2005). A common finding, replicated in studies on many different species, is that when agents are trained in a larger arena, feature cues have more of an effect than when agents learn the reorientation task in a small arena. While such a finding is common, it is not universal. For instance, although 3-year-old children can use wall color as a feature to reorient in a large room but not in a small room (Learmonth et al., 2002), the use of feature cues in the large room does not occur if movements of 3-year children are restricted in moving beyond the area covered by the small room (Learmonth et al., 2008).

Does arena size affect perceptron performance on the reorientation task? To answer this question, we simulated an experiment in which agents were trained in an arena of one size, with three-featured panels at each corner, and were then placed in an arena of a different size, with the objects moved clockwise to the right, producing an affine transformation of the training arena (Vallortigara et al., 2005). However, in addition to these conditions that were studied by Vallortigara et al, we also studied perceptron responses to arenas to which an affine transformation had been applied, but which remained the same size.

8.1. Method

Perceptrons were trained in two different conditions. In one, they were trained in a small rectangular arena that had a unique, three-featured landmark at each of the arena's four corners. The feature values for the object at the near location were identical to those used in the training phase of Simulations 3 and 5. In the small rectangular arena, the short walls were four arbitrary units in length, and the long walls were eight arbitrary units in length; all wall colors were black. The properties that defined each of the four locations were represented to the networks using the Table 1 coding scheme. In the second condition, perceptrons were trained in a large rectangular arena. This training condition was identical to the small arena condition, with the exception that the short walls were six arbitrary units in length, and the long walls were 12 arbitrary units in length. That is, the large arena had the same proportions as the small arena, but was 1.5 times larger.

Ten different perceptrons served as subjects in each condition, and were trained until convergence using the same procedure that was employed in all of the previous simulations. After training was completed, the responses of each perceptron were observed in four different arenas. Two of these were the small and the large arenas that were used for training. The third was the small arena to which an affine transformation had been applied, which was accomplished by moving each landmark to the nearest location in a clockwise direction. The fourth was the large arena to which the same affine transformation had been applied.

8.2. *Results and discussion*

On average, perceptrons trained in the small arena converged to a solution after 254.3 epochs of training. Perceptrons trained in the large arena converged slightly faster, requiring an average of 244.2 epochs of training.

8.2.1. *Perceptron responses*

The average responses of perceptrons to each of the test arenas are provided in the upper part of Table 7.

Even though the perceptron coding scheme was not designed with size effects in mind, and examination of Table 7 reveals that perceptrons trained in the small arena generated different responses than did perceptrons trained in the large arena. Some of these differences are analogous to those that have been reported in the literature. For instance, when perceptrons that have been trained in the large arena are tested in the small arena with moved landmarks, they have a stronger preference for the near location (which is marked by the previously reinforced set of feature cues) than is the case for perceptrons that have been trained in the small arena and are tested in the affine transformation of the large arena. Studies with human subjects have shown that those trained in a large arena are more sensitive to features when tested in a small arena than are subjects who have been trained in a small arena and are tested in a large arena (Ratliff & Newcombe, 2008); similar results have been demonstrated in studies of fish (Sovrano, Bisazza, & Vallortigara, 2005).

Another general property of these results is that perceptrons solved the reorientation problem by exploiting both geometric and feature cues, a result that has been found in previous studies of size effects (Chiandetti et al., 2007; Vallortigara et al., 2005). For instance, note that perceptrons trained in the small arena generated moderately strong responses (0.28) to the correct and rotational locations in the small arena that has undergone an affine transformation. These two locations are geometrically equivalent, and this geometry was reinforced during training. Furthermore, this sensitivity to geometric cues is transferred when perceptrons trained in the small arena are tested in the large arena that has undergone an affine transformation. In this case, the responses to the rotational and correct locations are slightly weaker (0.16) than in the small arena, but are still substantially greater than to the geometrically incorrect far location. A similar pattern of re-

Table 7
Average responses of perceptrons to landmarks in different sized arenas after training in Simulation 6.

Training condition	Arena location	Testing arena			
		Small arena		Large arena	
		Landmarks in original locations	Landmarks moved in clockwise direction	Landmarks in original locations	Landmarks moved in clockwise direction
Perceptron activity					
Small arena	Near	0.04	0.46	0.05	0.57
	Rotational	0.09	0.28	0.05	0.16
	Far	0.04	0.01	0.05	0.01
	Correct	0.90	0.28	0.81	0.16
Large arena	Near	0.10	0.67	0.03	0.37
	Rotational	0.06	0.22	0.09	0.32
	Far	0.10	0.02	0.03	0.01
	Correct	0.84	0.22	0.90	0.33
Choice rule					
Small arena	Near	0.04	0.45	0.05	0.63
	Rotational	0.08	0.27	0.05	0.18
	Far	0.04	0.01	0.05	0.01
	Correct	0.84	0.27	0.84	0.18
Large arena	Near	0.09	0.59	0.03	0.36
	Rotational	0.05	0.19	0.09	0.31
	Far	0.09	0.02	0.03	0.01
	Correct	0.76	0.19	0.86	0.32

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

sponses is observed for the perceptrons trained in the large arena.

However, other size-related differences do not reflect extant experimental results. For instance, from Table 7 it would appear that when perceptrons are tested in an affine transformation of the same-sized arena in which they were trained, perceptrons trained in the large arena have a lower response to the previously reinforced feature cues (now in the near location) than is the case for perceptrons trained in the small arena. The opposite effect has been reported in a study of spatial reorientation by chicks (Vallortigara et al., 2005). As well, when placed in the same-sized arena that has undergone an affine transformation, the perceptrons trained in the large arena have a higher response to the correct geometric cues than do perceptrons trained in the small arena. Again, these results are inconsistent with the more common findings that training in large arenas produces higher sensitivity to feature cues, and lower sensitivity to geometric cues, than does training in small arenas.

8.2.2. Choice probabilities

The bottom half of Table 7 presents the results of this simulation after perceptron responses have been converted into choice probabilities. This version of the results preserves the general interpretations that were presented when perceptron responses were discussed: choice probabilities are affected by the size of the arena in which training occurred, both geometric and feature cues are processed, and some of the size effects mirror those found in the literature while others do not. One interesting discrepancy between the lower and upper half of Table 7 concerns the choice probability to the near location when features have been moved. When choice probabilities are computed, the perceptron trained in the small arena now has a slightly higher choice probability to the near location in the large test arena than does the perceptron trained in the large arena to the near location in the small test arena. This result is opposite to that seen when perceptron responses are examined, and is inconsistent with the general finding in the literature that feature cues are more salient to agents trained in large arenas than to agents trained in smaller arenas. Again, this discrepancy demonstrates the need to report multiple dependent measures in studies of the reorientation task, whether they are simulations or experimental studies of biological agents.

To summarize, the responses of perceptrons trained to solve the reorientation task are sensitive to arena size, at least when wall lengths are encoded with the thermometer code detailed in Table 1. Some of the size effects on perceptron responses resemble those that have been observed in previous experimental studies, but others do not. Given the complex nature of size effects in the literature (Learmonth et al., 2008), it is difficult to know what to make of these inconsistencies. It may be the case that size effects reflect complexities of the organism (such as movement, genetic predisposition, or prior experience; see Learmonth et al., 2008) rather than inherent aspects of the geometric and featural stimuli in the environment. For example, biological factors such as movement, development, and predisposition may change the way in which geometric and featural information is encoded or weighed. At the very least, perceptron sensitivity to arena size indicate that perceptrons can provide an interesting medium to explore a variety of hypotheses about arena size and reorientation.

9. General discussion

The current paper explored how a particularly simple type of artificial neural network, the perceptron (Rosenblatt, 1958, 1962) could accomplish spatial reorientation. The simulations reported above indicate that perceptrons can generate some of the basic findings associated with the reorientation task (Cheng, 2005; Cheng & Newcombe, 2005): rotational error in rectangular arenas that do not include feature cues, the ability to uniquely identify target locations when featural cues are present, and the use of both feature and geometric cues when both are available, but when only one type of cue (i.e., feature) is actually required. The perceptrons were able to generate these effects in arenas of different shapes and sizes, and using feature cues of varying types. Importantly, interesting quantitative differences were revealed when results were reported in terms of perceptron responses to cues at individual locations were converted into choice probabilities. This suggests that future studies of the reorientation task – be they simulations or experimental studies of biological agents – might need to report multiple dependent measures, because choice probabilities (the usual dependent measure) may not necessarily reflect underlying processes used to solve the reorientation task.

The fact that an extremely simple associative model can produce reorientation task regularities is important given the current status of the assumption that geometric processing in the reorientation task is modular (Cheng, 1986; Gallistel, 1990). Questions that have been raised about this strong assumption have recently led to the consideration of alternative notions of modularity (e.g., Cheng, 2005; Cheng & Newcombe, 2005), and have also led to models in which the assumption is abandoned altogether (Cheng, 2008; Cheung et al., 2008; Miller & Shettleworth, 2007, 2008; Stuerzl et al., 2008). The current simulations used an extremely simple encoding that was nonmodular, and in which geometric and feature information was represented as locally available cues that were not qualitatively different. They revealed – perhaps to the surprise of some readers – that these local, nonmodular codes still permitted the perceptrons to demonstrate that geometric cues were processed, even in situations in which available feature cues made such processing redundant. Such data has typically been used to argue for geometric modularity. However, we were able to simulate such results without appealing to modularity. It would appear that the general associative mechanisms embodied in the perceptron at the very least can provide a plausible account of spatial reorientation.

Of course, the perceptrons that have been described in the current paper are not the only models that have been applied to the reorientation task. Let us briefly consider the relationship of the perceptrons to other models that have appeared in the literature.

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

The model most similar to the perceptron is a modification of the Rescorla–Wagner learning rule to produce an operant learning paradigm for the reorientation task (Miller, 2009; Miller & Shettleworth, 2007, 2008). These two approaches are similar in that both models deal with the reorientation task by applying a general learning rule to geometric and featural cues Furthermore, it has long been known that the error-correcting rules used to train perceptrons can be translated into the Rescorla–Wagner model (Gluck & Bower, 1988; Gluck & Myers, 2001; Sutton & Barto, 1981). However, even with these similarities there are important differences between the two models.

First, in spite of the formal equivalence between perceptron learning and the Rescorla–Wagner formalism, it is still possible to generate behavioral differences between perceptrons and the Rescorla–Wagner model (Dawson, 2008; Dawson & Spetch, 2005) because of different assumptions made when the two different theories of learning are translated into working algorithms. We saw earlier in Simulation 5 that this holds true for the reorientation task, because the Miller/Shettleworth model generated different results than did the perceptrons. In particular, the perceptrons were more influenced by the twin processes of excitation and release from inhibition.

Second, Miller and Shettleworth (2007) emphasize the operant learning context of the reorientation task. The basic idea is that as an agent is reinforced for exploring a particular location, it will be more likely to explore that location, and less likely to explore others. They modify the Rescorla–Wagner model to reflect the changing probabilities of visiting the different locations. They also argue that the operant nature of learning during reorientation is responsible for an effect that they call feature enhancement, in which certain cues acquire greater associative strength than would be expected. Feature enhancement is evident, for example, when learning about some cues is not overshadowed by learning about others – for instance, when geometric information is utilized even when feature information alone would suffice. In contrast, the perceptrons trained in the current manuscript do not employ operant procedures, but still demonstrate feature enhancement effects. Indeed, we have developed an operant training procedure for perceptrons, but have found that it does not lead to different final networks than does the more typical nonoperant training (Dawson et al., 2009).

A third difference between the two models concerns the ease with which the two can be extended when necessary. Miller and Shettleworth (2007) point out that in some instances their model makes predictions that are contrary to the results of extant experiments. The same is true for the current version of the perceptron models, as shown in the study of size effects in Simulation 6. How might future models be developed to deal with this problem? On the one hand, it is not immediately obvious what would be done with a Rescorla–Wagner based model such as the one proposed by Miller and Shettleworth. On the other hand, it is well-known that the perceptron is a simple artificial neural network that has limitations in the regularities that it can detect, and that one can overcome such limitations by moving to a multilayer perceptron that includes hidden units, and is trained by a variant of the error-correcting rule described above. Thus, if required, the type of work described in the paper above could be extended by training multilayer perceptrons on the reorientation task. However, given the successes noted above, the need for this extension is not an immediate one.

The perceptron model is less obviously related to some others that have been applied to the reorientation task, but is not incompatible with them, and may complement them. One is the view-based model in which agents reorient themselves to a new arena by moving towards a location that is most similar to remembered views (Cheung et al., 2008; Stuerzl et al., 2008). The features involved in this approach can be very primitive (e.g., pixel-based representations that have undergone little image analysis). This theory can predict phenomena like rotational error, but does not yet include a component that explains the mechanisms by which particular views are remembered when reinforced. It is possible that artificial neural networks, whose inputs represent information currently used by view-based theories instead of the highly processed features described in our simulations, could provide a learning mechanism for this theory. Similarly, theories have been proposed that explain reorientation by appealing to the adaptive weighting of the validity of various cue features (Cheng et al., 2007; Newcombe & Ratliff, 2007). These theories have described this adaptive combination in general terms, and have not provided detailed algorithms for how features would be weighted and combined as a function of experience. Neural networks like those described in this paper might provide such an algorithm. For instance, perceptrons are very adept at matching the probability of reinforcement associated with different cues (Dawson et al., 2009), and there are strong formal links between neural networks and probabilistic accounts of the relationship between cues and outcomes such as contingency theory (Shanks, 1995).

In summary, the perceptrons described in this paper have been shown to generate important known results in the reorientation task, and are flexible enough to generate novel predictions to be tested in new experiments. The behavior of the perceptrons raises the possibility that the reorientation task can be solved by the associative mechanisms that are instantiated in artificial neural networks. This is important given that researchers who have championed geometric modularity are now revisiting some of their earlier proposals (e.g., Cheng, 2005). Future research is required to completely determine the successes and the failures of the particular artificial neural network described above, and to establish whether those interested in this particular navigation task may require more complicated multilayer networks in the future. Another interesting direction for future research that is afforded by using a neural network model will be to explore the relationship between the perceptron results and the neuroscience literature on navigation, given the fact that artificial neural networks are intended to ultimately link behavior to brain (Enquist & Ghirlanda, 2005), as well as the fact that there is a growing interest in the neuroscience of spatial cognition (Burgess, 2008).

References

Batty, E. R., Bloomfield, L. L., Spetch, M. L., & Sturdy, C. B. (2009). Comparing black-capped (*Poecile atricapillus*) and mountain chickadees (*Poecile gambeli*): Use of geometric and featural information in a spatial orientation task. *Animal Cognition, 12*(4), 633–641.

Bechtel, W., & Abrahamsen, A. A. (2002). *Connectionism and the mind: Parallel processing, dynamics, and evolution in networks* (2nd ed.). Malden, MA: Blackwell.

Braitenberg, V. (1984). *Vehicles: Explorations in synthetic psychology*. Cambridge, MA: MIT Press.

Brooks, R. A. (1999). *Cambrian intelligence. The early history of the new AI*. Cambridge, MA: MIT Press.

Brown, A. A., Spetch, M. L., & Hurd, P. L. (2007). Growing in circles – rearing environment alters spatial navigation in fish. *Psychological Science, 18*(7), 569–573.

Burgess, N. (2008). Spatial cognition and the brain. *Year in Cognitive Neuroscience, 2008*(1124), 77–97.

Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition, 23*, 149–178.

Cheng, K. (2005). Reflections on geometry and navigation. *Connection Science, 17*(1–2), 5–21.

Cheng, K. (2008). Whither geometry? Troubles of the geometric module. *Trends in Cognitive Sciences, 12*(9), 355–361.

Cheng, K., & Newcombe, N. S. (2005). Is there a geometric module for spatial orientation? Squaring theory and evidence. *Psychonomic Bulletin and Review, 12*(1), 1–23.

Cheng, K., Shettleworth, S. J., Huttenlocher, J., & Rieser, J. J. (2007). Bayesian integration of spatial information. *Psychological Bulletin, 133*(4), 625–637.

Cheung, A., Stuerzl, 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), 15–30.

Chiandetti, C., Regolin, L., Sovrano, V. d., & Vallortigara, G. (2007). Spatial reorientation: The effects of space size on the encoding of landmark and geometry information. *Animal Cognition, 10*(2), 159–168.

Clark, A. (1997). *Being there: putting brain, body, and world together again*. Cambridge, Mass.: MIT Press.

Clark, A. (2003). *Natural-born Cyborgs*. Oxford; New York: Oxford University Press.

Dawson, M. R. W. (1998). *Understanding cognitive science*. Oxford, UK: Blackwell.

Dawson, M. R. W. (2004). *Minds and machines: Connectionism and psychological modeling*. Malden, MA: Blackwell Pub..

Dawson, M. R. W. (2005). *Connectionism: A hands-on approach* (1st ed.). Oxford, UK; Malden, MA: Blackwell Pub..

Dawson, M. R. W. (2008). Connectionism and classical conditioning. *Comparative Cognition and Behavior Reviews, 3*(Monograph), 1–115.

Dawson, M. R. W., Dupuis, B., Spetch, M. L., & Kelly, D. M. (2009). Simple artificial networks that match probability and exploit and explore when confronting a multiarmed bandit. *Ieee Transactions on Neural Networks, 20*(8), 1368–1371.

Dawson, M. R. W., & Spetch, M. L. (2005). Traditional perceptrons do not produce the overexpectation effect. *Neural Information Processing – Letters and Reviews, 7*(1), 11–17.

Dawson, M. R. W., & Zimmerman, C. (2003). Interpreting the internal structure of a connectionist model of the balance scale task. *Brain and Mind, 4*, 129–149.

Enquist, M., & Ghirlanda, S. (2005). *Neural networks and animal behavior*. Princeton: Princeton University Press.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning – An adaptive network model. *Journal of Experimental Psychology-General, 117*(3), 227–247.

Gluck, M. A., & Myers, C. (2001). *Gateway to memory: an introduction to neural network modeling of the hippocampus and learning*. Cambridge, Mass.: MIT Press.

Graham, M., Good, M. A., McGregor, A., & Pearce, J. M. (2006). Spatial learning based on the shape of the environment is influenced by properties of the objects forming the shape. *Journal of Experimental Psychology-Animal Behavior Processes, 32*(1), 44–59.

Gray, E. R., Bloomfield, L. L., Ferrey, A., Spetch, M. L., & Sturdy, C. B. (2005). Spatial encoding in mountain chickadees: features overshadow geometry. *Biology Letters, 1*(3), 314–317.

Gulliksen, H. (1953). A generalization of Thurstone's learning function. *Psychometrika, 18*(4), 297–307.

Hayward, A., Good, M. A., & Pearce, J. M. (2004). Failure of a landmark to restrict spatial learning based on the shape of the environment. *Quarterly Journal of Experimental Psychology Section B-Comparative and Physiological Psychology, 57*(4), 289–314.

Hermer, L., & Spelke, E. S. (1994). A geometric process for spatial reorientation in young children. *Nature, 370*(6484), 57–59.

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior, 13*(2), 243–266.

Kelly, D. M., Spetch, M. L., & Heth, C. D. (1998). Pigeons' (Columba livia) encoding of geometric and featural properties of a spatial environment. *Journal of Comparative Psychology, 112*(3), 259–269.

Learmonth, A. E., Nadel, L., & Newcombe, N. S. (2002). Children's use of landmarks: Implications for modularity theory. *Psychological Science, 13*(4), 337–341.

Learmonth, A. E., Newcombe, N. S., & Huttenlocher, J. (2001). Toddlers' use of metric information and landmarks to reorient. *Journal of Experimental Child Psychology, 80*(3), 225–244.

Learmonth, A. E., Newcombe, N. S., Sheridan, N., & Jones, M. (2008). Why size counts: children's spatial reorientation in large and small enclosures. *Developmental Science, 11*(3), 414–426.

Luce, R. D. (1959). *Individual choice behavior*. New York: Wiley.

Luce, R. D. (1961). A choice theory analysis of similarity judgments. *Psychometrika, 26*(2), 151–163.

Luce, R. D. (1977). The choice axiom after 20 years. *Journal of Mathematical Psychology, 15*(3), 215–233.

Miller, N. Y. (2009). Modeling the effects of enclosure size on geometry learning. *Behavioural Processes, 80*, 306–313.

Miller, N. Y., & Shettleworth, S. J. (2007). Learning about environmental geometry: An associative model. *Journal of Experimental Psychology: Animal Behavior Processes, 33*, 191–212.

Miller, 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), 419–422.

Newcombe, N. S., & Ratliff, K. R. (2007). Explaining the development of spatial reorientation: Modularity-plus-language versus the emergence of adaptive combination. In J. Plumert & J. Spencer (Eds.), *The emerging spatial mind* (pp. 53–76). Oxford: Oxford University Press.

Newell, A. (1973). Production systems: Models of control structures. In W. G. Chase (Ed.), *Visual information processing* (pp. 463–526). New York, NY: Academic Press.

Newell, A. (1980). Physical symbol systems. *Cognitive Science, 4*, 135–183.

Nolfi, S. (2002). Power and limits of reactive agents. *Neurocomputing, 42*, 119–145.

Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics*. Cambridge, Mass.: MIT Press.

Pearce, J. M. (1997). *Animal learning and cognition: an introduction*. East Sussex: Psychology Press.

Pearce, J. M. (2002). Evaluation and development of a connectionist theory of configural learning. *Animal Learning and Behavior, 30*(2), 73–95.

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), 135–147.

Pearce, J. M., Graham, M., Good, M. A., Jones, P. M., & McGregor, A. (2006). Potentiation, overshadowing, and blocking of spatial learning based on-the shape of the environment. *Journal of Experimental Psychology-Animal Behavior Processes, 32*(3), 201–214.

Pearce, J. M., Ward-Robinson, J., Good, M., Fussell, C., & Aydin, A. (2001). Influence of a beacon on spatial learning based on the shape of the test environment. *Journal of Experimental Psychology-Animal Behavior Processes, 27*(4), 329–344.

Pfeifer, R., & Scheier, C. (1999). *Understanding Intelligence*. Cambridge, MA: MIT Press.

Pylyshyn, Z. W. (1984). *Computation and cognition*. Cambridge, MA: MIT Press.

Ratliff, K. R., & Newcombe, N. S. (2008). Reorienting when cues conflict: Evidence for an adaptive-combination view. *Psychological Science, 19*(12), 1301–1307.

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.

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review, 65*(6), 386–408.

Rosenblatt, F. (1962). *Principles of neurodynamics*. Washington: Spartan Books.

Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing, V.1*. Cambridge, MA: MIT Press.

Shanks, D. R. (1995). *The psychology of associative learning*. Cambridge, UK: Cambridge University Press.

Sovrano, V. A., Bisazza, A., & Vallortigara, G. (2002). Modularity and spatial reorientation in a simple mind: encoding of geometric and nongeometric properties of a spatial environment by fish. *Cognition, 85*(2), B51–B59.

Sovrano, V. A., Bisazza, A., & Vallortigara, G. (2003). Modularity as a fish (*Xenotoca eiseni*) views it: Conjoining geometric and nongeometric information for spatial reorientation. *Journal of Experimental Psychology-Animal Behavior Processes, 29*(3), 199–210.

Sovrano, V. A., Bisazza, A., & Vallortigara, G. (2005). Animals' use of landmarks and metric information to reorient: effects of the size of the experimental space. *Cognition, 97*(2), 121–133.

Sovrano, V. A., Bisazza, A., & Vallortigara, G. (2007). How fish do geometry in large and in small spaces. *Animal Cognition, 10*(1), 47–54.

Sovrano, V. A., & Vallortigara, G. (2006). Dissecting the geometric module: A sense linkage for metric and landmark information in animals' spatial reorientation. *Psychological Science, 17*(7), 616–621.

Stuerzl, W., Cheung, A., Cheng, K., & Zeil, J. (2008). The information content of panoramic images I: The rotational errors and the similarity of views in rectangular experimental arenas. *Journal of Experimental Psychology-Animal Behavior Processes, 34*(1), 1–14.

Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review, 88*(2), 135–170.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning*. Cambridge, MA: MIT Press.

Thurstone, L. L. (1930). The learning function. *Journal of General Psychology, 3*, 469–493.

Vallortigara, G., Feruglio, M., & Sovrano, V. A. (2005). Reorientation by geometric and landmark information in environments of different size. *Developmental Science, 8*(5), 393–401.

Vallortigara, G., Zanforlin, M., & Pasti, G. (1990). Geometric modules in animals spatial representations: A test with chicks (*Gallus-Gallus-Domesticus*). *Journal of Comparative Psychology, 104*(3), 248–254.

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 and Behavior, 32*(3), 289–298.

Wystrach, A., & Beugnon, G. (2009). Ants learn geometry and features. *Current Biology, 19*(1), 61–66.

Please cite this article in press as: Dawson, M. R. W., et al. Using perceptrons to explore the reorientation task. *Cognition* (2009), doi:10.1016/j.cognition.2009.09.006