The Information Content of Panoramic Images II: View-Based Navigation in Nonrectangular Experimental Arenas

Allen Cheung, Wolfgang Stürzl, and Jochen Zeil
The Australian National University

Ken Cheng
Macquarie University

Two recent studies testing navigation of rats in swimming pools have posed problems for any account of the use of purely geometric properties of space in navigation (M. Graham, M. A. Good, A. McGregor, & J. M. Pearce, 2006; J. M. Pearce, M. A. Good, P. M. Jones, & A. McGregor, 2004). The authors simulated 1 experiment from each study in a virtual reality environment to test whether experimental results could be explained by view-based navigation. The authors recorded a reference image at the target location and then determined global panoramic image differences between this image and images taken at regularly spaced locations throughout the arena. A formal model, in which an agent attempts to minimize image differences between the reference image and current views, generated trajectories that could be compared with the search performance of rats. For both experiments, this model mimics many aspects of rat behavior. View-based navigation provides a sufficient and parsimonious explanation for a range of navigational behaviors of rats under these experimental conditions.

Keywords: spatial learning, view-based navigation, computer simulation, rats, geometry

Supplemental materials: http://dx.doi.org/10.1037/0097-7403.34.1.15.supp

The spatial abilities of laboratory rats often have been tested in small experimental spaces, such as a radial maze with arms radiating from a central platform (Olton & Samuelson, 1976; Roberts, 1984), a swimming pool (Morris, 1981), or an enclosed rectangular box (Cheng, 1986; reviewed by Cheng & Newcombe, 2005). Much of this research has been concerned with the neurobiological foundations of spatial behavior (Jeffery, 2003). The performance of rats in rectangular boxes has led to the hypothesis that they remember and are guided by the geometric shape of their environment (Cheng, 1986; reviewed by Cheng & Newcombe, 2005), an ability possibly based on a dedicated module for encoding geometry (Wang & Spelke, 2002, 2003; see, however, Newcombe, 2002).

The impetus for the idea of a geometric module came originally from the observation that rats searching for food in one of the corners of an enclosed rectangular arena made characteristic and systematic errors (Cheng, 1986). The rats often confused diagonally opposite corners, making what has come to be called a *rotational error*. They made rotational errors even when the inside of the rectangular arena provided distinct and unique cues, such as when one wall was white and three walls were black, or when different patterns identified each of the corners. It thus appeared that rats ignored all these internal cues and paid more attention to the overall shape of the box.

A number of proposals were made to account for this peculiar search behavior (Cheng, 1986, 2005; Cheng & Gallistel, 2005; Gallistel, 1990; Pearce, Good, Jones, & McGregor, 2004; Tommasi & Polli, 2004). Pearce et al. (2004), for instance, trained rats in a rectangular pool and transferred them to a kite-shaped pool (see inset in Figure 1). This test pool had a different shape from the training pool, and matching on the basis of overall shape was thus impossible. On the basis of the fact that the rats picked two of the corners in the kite preferentially, the authors suggested that rats did not encode the global shape of the space but used local geometric properties. Tommasi and Polli (2004) found nonrandom choices in chicks given transformation tests in shapes different from the parallelogram shape used in training, and they concluded likewise. Cheng and Gallistel (2005) and Cheng (2005) suggested in turn that the use of some global shape parameters, such as principal axes of space, can account for the data just as well. McGregor, Jones, Good, and Pearce (2006), however, provided experimental evidence that rats do not use the principal axis. All of these proposals have in common the assumption that to solve the task of relocating a reward in a corner, rats have access to explicit infor-
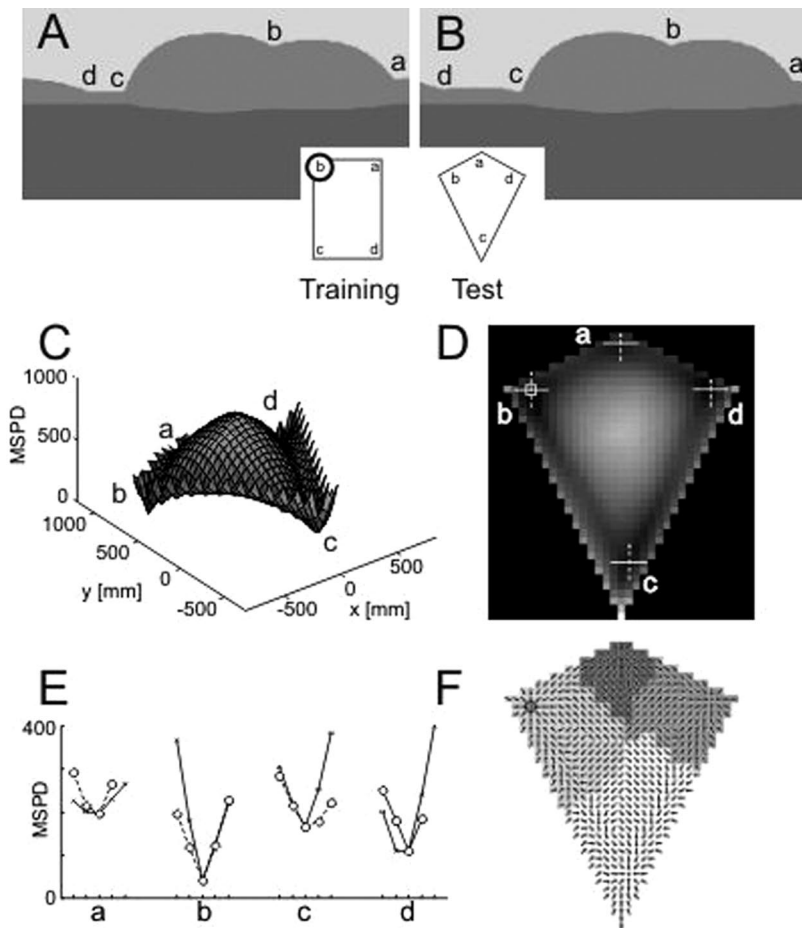
Figure 1. Experimental setup by Pearce et al. (2004): Rats were trained to find a hidden platform close to one of the corners of a rectangular swimming pool with four white walls (inset A). The rats were then tested in a kite-shaped pool (inset B). (A) The panoramic view of the training corner b (see inset) as simulated in a virtual rectangular box. The panoramic view wraps around so that the right edge of the image coincides with the left edge. (B) The panoramic view of the equivalent corner b in the kite-shaped test box. (C) The three-dimensional representation of the image difference function (IDF) of the reference image in corner b of the rectangular training box in the kite-shaped box. The *z*-axis (vertical) measures the mean squared pixel difference (MSPD) between the image at that point in the arena and the reference image. The saw-tooth edges of the IDF are a simulation artifact caused by the coarse sampling grid. (D) The two-dimensional representation of the same IDF, with gray-level code indicating the local image difference values (high–white; low–black). The location of the "target" corner, which corresponds to the training corner b in the rectangular arena, is marked by a square. Crossed continuous and dashed lines indicate the direction of transects through the mimina of the IDF as shown in E. (E) Local transects through the local minima of the IDF close to the corners of the kite-shaped test arena. (F) The local slopes of the IDF (arrows) and the catchment areas of local minima (gray-shaded areas).

mation about the geometric properties of the environment, such as the local geometry at that corner (Pearce et al., 2004; Tommasi & Polli, 2004) or global parameters derived from shape information (Cheng, 2005; Cheng & Gallistel, 2005; Gallistel, 1990).

Some recent results, however, have posed grave problems for any theory claiming that only geometric properties independent of features are extracted and used. Graham, Good, McGregor, and Pearce (2006) trained and tested rats in a kite-shaped swimming pool. Unlike a rectangle, this space has no geometric ambiguities if the overall shape of the arena is considered. In one training condition supposedly conducive to the use of geometric informa-

tion, the target corner (where a submerged escape platform was located) was always at one of the right-angled corners, but the colors of the walls varied from trial to trial. The target corner was formed by white walls on some trials and by black walls on other trials. Likewise, the walls opposite the target also varied between black and white. Under these training conditions, the colors of walls, a major featural property, were uninformative, whereas geometric information was unambiguous for locating the target. If the extraction of geometric information is "obligatory," as Cheng and Newcombe (2005) and Cheng (2005) have suggested, rats should learn the task readily. In fact, they found it very difficult.

However, they found it much easier if the walls adjacent to the platform were of a different color from those adjacent to the opposite right-angled corner. Such results suggest that perhaps another explanation is required.

In the companion article (Stürzl, Cheung, Cheng, & Zeil, 2007), we showed that geometric information is contained implicitly in panoramic images taken from the viewpoint of animals in such experimental spaces. We reconstructed views in the experimental environment first used by Cheng (1986) and found that an agent attending only to image differences would make rotational errors, without any explicit computation of geometric properties, whether they be local or global. We showed this by determining the catchment areas of each local minimum in image difference between the reference image and the image at that location. The reference image is the panoramic image taken at the target location. A catchment area is the region of space from which an agent would end up at a local (or global) minimum in image difference (see also Cartwright & Collett, 1987; Zeil, Hofmann, & Chahl, 2003). We assumed that the agent followed a local gradient-descent algorithm based on image differences.

The most important result of this study by Stürzl et al. (2007) was that under a variety of conditions, local minima with substantial catchment areas were found not only at the training or target corner, but also at the corner diagonally opposite the training location, the location of the rotational error in animal experiments. It is important to note that these rotational error catchments were found even when one wall was white and three walls were black, like in some of Cheng's (1986) original experiments.

In the companion article, we also showed that experiments in a virtual reality environment are fully compatible with experiments done with a real imaging device and produce qualitatively similar results (Stürzl et al., 2007). In the present study, we tested whether view-based navigation can also explain the performance of rats placed in a kite-shaped swimming pool as reported by Pearce et al. (2004) and Graham et al. (2006). These experiments were chosen because they have posed problems for the notion of geometry-based navigation. Data from Pearce et al. pose grave problems for the hypothesis that global shape information is used. Data from Graham et al. showed that rats had great difficulties in using geometric information of any kind under conditions when it would have benefitted them to do so. We therefore reconstructed the training and test spaces used by Pearce et al. and Graham et al. in a virtual environment and simulated the performance of a view-based homing strategy. As in our previous work (Stürzl et al., 2007), a reference image was recorded at the target corners to which rats were trained, and we then calculated the differences between these reference images and images over the entire space of the virtual pool. Local minima were identified, and their catchment areas were calculated using a local gradient-descent process. Finally, a formal performance model incorporating all these processes was created.

Experiment 1: Simulations of Pearce et al. (2004)

We first simulated the conditions of Experiment 1 of Pearce et al. (2004) in which rats were trained in a white rectangular pool and transferred to a white kite-shaped pool for testing (Experiment 1A) or vice versa (Experiment 1B).

Method

The virtual arena. The details of the construction of experimental arenas in virtual reality are described in the companion article (Stürzl et al., 2007). Briefly, to construct images of experimental arenas, texture bitmaps were created in MATLAB (Mathworks, Natick, MA) for all surfaces, including the "ground" and the "sky," using desired gray pixel values from 0 to 255, corresponding to camera settings of gain 2.25 and gamma off (see Stürzl et al., 2007). The textures were single homogeneous gray values for each entire rectangular surface. Relevant shapes (e.g., rectangles) were drawn using OpenGL, and the corresponding textures were pasted onto these shapes. No light source was added so that there would be no variation of pixel values depending on the assumed direction of illumination. In the real arenas, the light source was from above, as was the case in all of the experiments simulated, making the ceiling the brightest part of the virtual environment. In the absence of relevant image data from the original experimental arenas, we substituted pixel values from real images from within black and white rectangular boxes from Stürzl et al. (2007). The actual gray values used are documented in Table 1. The sizes of the virtual arenas were the same as the real arenas used by Pearce et al. (2004): The long walls of both the rectangular and kite-shaped arena were 1.8 m long and the short walls 0.9 m long.

The details of constructing panoramic views of this environment are given in Stürzl et al. (2007). Briefly, rectangular panoramic images representing 360° in azimuth and 180° in elevation were created by a custom-made program (Allen Cheung, The Australian National University), which takes position and orientation values as input and returns the image as its output. In the current simulations, each arena was sampled at every 5 cm in x- and y-directions. The target corner in all simulations was at a right-angled corner, 10 cm from the long wall and 20 cm from the short wall, viewed at 2 cm above the surface (which was the water surface in the experiments in question).

The image difference function (IDF). The image difference between two images was calculated as the mean squared pixel difference (MSPD) over all corresponding pairs of pixels in the two images (for details, see Stürzl et al., 2007). To compare two panoramic images, they must be "lined up" with one another. Each view can be imagined as being pasted onto the inside of a cylinder, and the two cylinders need to be oriented correctly with respect to one another. Because we made no a priori assumptions about an animal's heading direction during the various navigation tasks, we computed the MSPD for all possible headings in steps of 1°, and then we used the lowest MSPD at that particular location. The IDF is the two-dimensional distribution of MSPD values between a reference image and the images taken at each location inside the arena (for details, see Cartwright & Collett, 1987; Stürzl et al., 2007; Zeil et al., 2003).

Both the method of view-based homing and the definition of a catchment area have been described in detail elsewhere (Stürzl et al., 2007; Zeil et al., 2003). Briefly, image differences relative to the reference image were calculated within a grid of 24 neighboring and next-but-one units (with 5-cm spacing) around every position in the virtual space to determine the local slope of the IDF at each of these locations. These local slopes thus map out the topography of the IDF, indicating how a simulated agent would

Table 1
Dimensions and Grayscale Values of the Virtual Arenas Used in the Simulations

Arena	Wall length (cm)		Wall height (cm)	Pixel value			
	1	2		"Sky"	"Water"	Light wall	Dark wall
Pearce et al. (2004)							
Kite (W)	180	90	34	214	111	138	—
Rectangle (W)	180	90	34	214	111	138	—
Kite (B&W)	180	90	34	214	111	138	2
Graham et al. (2006)							
Kite (B&W)	180	90	34	204	111	177	48
Kite (W)	180	90	34	204	111	177	—
Square (B&W)	141	141	34	204	111	177	48

Note. B = black; W = white. Pixel values are integers between 0 and 255 (8-bit integers). The pixel values represent the final output used for calculating mean squared pixel differences (MSPD) values, after the unwarping transformation. The discrepancy between initial texture gray value and the final output gray value differed by a maximum of 1.2 units in these simulations (mostly attributable to rounding errors).

move at each location if it followed the direction of the lowest MSPD value in the local neighborhood. If the IDF is smooth, such a gradient-descent algorithm would end up in those locations at which the IDF has a local or an absolute minimum (global minimum). Each minimum of the IDF has a catchment area: an area in the virtual arena from within which a simulated gradient-descending agent would end up at that particular minimum (for details, see Stürzl et al., 2007; Zeil et al., 2003). It should be emphasized that during gradient descent, an animal does not need to know its allocentric location but merely finds the correct heading. Details are provided in the *Results and Discussion* section.

Modeling performance. We modeled the performance of an agent searching for a goal in the experimental arenas in the following way (for details of the implementation, see Appendix in the online supplementary materials). The local IDF gradient dictates a heading direction purely on the basis of local IDF values, which in the current form of the model is assumed to be known. In practice, the local IDF gradient can be estimated by a local search of the IDF to find the minimum value within some predefined search radius (see Stürzl et al., 2007; Zeil et al., 2003). We added white noise (completely independent of IDF values) to this local heading direction. Noise is a biological reality, but the addition of noise also allows the agent to escape from shallow false/local minima, especially when the other side of such a minimum is a barrier (in our case, a wall).

Finally, a heading filter was applied to simulate the fact that for rats in water, in the absence of arena information, the best moment-to-moment prediction of future heading is likely to be current heading. In other words, we assumed that there is resistance against changing heading direction because of physical limitations, including actual inertia, biomechanics of swimming, or energetic constraints, as well as possible navigational reasons (e.g., to facilitate the interpretation of view transformations).

With these three main components in the performance model—the use of the IDF gradient, the addition of noise, and the application of a heading filter—we could simulate individual paths for a given IDF, the amount of noise, and the width of the heading filter. For each experiment, we carried out at least 1,000 simula-

tions. Agents were released with equal probability facing each of the four walls, as Pearce et al. (2004) did with rats. Simulated release points were distributed in a Gaussian fashion with a radial standard deviation of 10 cm (see inset in Figure 4, which will be discussed in the *Results and Discussion* section), and release heading had a mean direction facing the wall with a Gaussian error distribution with 11.5° standard deviation, independent of release position. Because the platform was removed in test trials and because experiments often had time limits, simulated agents were allowed to take 200 steps of 5-cm step length (i.e., a total path length of 10 m) before simulation was stopped.

Three performance models were compared with the performance of rats on the behavioral measure of the percentage of first entries into each of the corners of the enclosure, defined as a circular arc of 40-cm radius centered at the vertices of the arena (see Figure 4). We compared the performance models, which we called *pure noise*, *pure IDF*, and *optimal search*. In the pure noise condition, the movement directions of the agent were independent of the local gradient of the IDF and were determined by random noise that was added to a broad heading filter. In this condition, a direction was randomly selected at the release point, and the heading filter tended to keep the agent moving in that direction. In the pure IDF condition, the movement direction was determined by the local slope of the IDF and the same broad heading filter that led to a tendency to continue moving into the current direction of heading. In the optimal search condition, the direction was also determined by the local gradient of the IDF, but random noise was added before applying the heading filter. The optimal amplitude of the noise ($\alpha = 0.1$) and the optimal width of the heading filter ($\sigma = 0.4$ radians; see Appendix in the online supplementary materials) were determined by finding the values that produced the best performance of the model in mimicking the behavior of rats on two sets of tests conducted by Graham et al. (2006, their Figure 3). For details, see Experiment 2.

Results and Discussion

In their Experiment 1A, Pearce et al. (2004) trained rats to find a submerged platform in one of the corners of a white rectangular

pool (for instance, corner b; see insets in Figure 1) and then asked how they would transfer what they learned when forced to search for the platform in a kite-shaped pool. Their results showed that rats persistently searched for the platform at the equivalent corner b and at the apex c of the kite-shaped test space (see insets in Figure 1). We simulated these conditions in a virtual reality environment and explain our procedures and the presentation of results in Figure 1, with subsequent figures following the same conventions.

Figure 1A shows the panoramic image taken at the target corner of the rectangular training pool (see inset); the ceiling, where the light comes from, is the brightest feature in the image, followed by the walls and the dark "ground" (water). Figure 1B shows the best matching image in the kite-shaped test pool, at location b in the inset. It looks indeed similar to the reference image, but differences are clearly discernible in the distant part of the arena.

Figure 1C shows the IDF over the entire test pool. The z-axis (dependent measure) shows the MSPD in the pixel-by-pixel comparison between the image at location (x, y) in the pool and the reference image taken in the training pool. The saw-tooth edges of the IDF are due to unavoidable artifacts in the virtual reality environment because the regular sampling grid does not fit smoothly into a kite-shaped space and because locations very close to the walls of the arena suffer from artifacts in OpenGL, which cause surface elements at the edges to erratically appear or disappear, resulting in large MSPD changes. The IDF forms a hilltop at the center of the kite-shaped test space, where image differences are largest, with image differences decreasing smoothly toward the edges of the test space. The fact that the lowest image difference values are found in the corners of the test space is documented in Figure 1D, which shows a two-dimensional, gray-level coded representation of the IDF. These minima are actually deep troughs of MSPD values, as can be seen in Figure 1E, which shows transects along the x-direction (solid lines) and the y-direction (dashed lines) through the IDF at these locations. The locations and directions of these transects are indicated by the equivalent solid and dashed lines in Figure 1D; the absolute minimum is marked by a small square. It is not zero in this case because of the relatively coarse grid distances used in the simulation and because the global geometry of training and test space is different.

Pearce et al. (2004) found that rats trained to position b in a rectangular pool would predominantly search in the kite-shaped test pool at locations b and c. Our simulation shows that image differences within this test space have a global minimum at b, but a comparatively shallow local minimum near c. Although the location of the global minimum at b does coincide with the corner rats most frequently visit, the relative depth of minima on their own thus does not reflect the search behavior of rats, which prefer to visit corner c over the remaining corners. However, from their release points at the center of the four walls, rats navigating by image matching cannot compare the relative depths of image difference minima throughout the arena. Instead, they need to probe and follow the local gradient of image differences. A way of showing what local information they would have in this test situation is to map the catchment areas for this particular IDF (see *Method* section and Stürzl et al., 2007), which we have done in Figure 1F. The arrows in Figure 1F show the direction of the local slope of the IDF, and we marked the catchment areas over which these local slopes lead to a particular minimum with different

shades of gray. It is important to note that all of these different catchment areas are based on image differences relative to one single reference image, taken at position b in the rectangular training pool (see inset in Figure 1A). Clearly, the largest catchment areas are those leading to corner b (the "correct corner" according to Pearce et al.) and to corner c (which Pearce et al. called the apex), exactly the corners most rats first visited directly on first encountering this test in the kite-shaped pool (Figure 2, right panel in Pearce et al.).

Pearce et al. (2004) also trained rats in a white kite-shaped pool and tested them in a white rectangular space. The results of our simulation are shown in Figure 2. The IDF in the rectangular test arena with the image from the training location b in the kite-shaped arena shows distinct minima in all four corners (Figures 2A and 2B), with lowest values at the diagonally opposite corners b and d (Figure 2C). Rats were released facing the center of one of the four walls. We have drawn a semicircle with a radius of 20 cm around these locations in Figure 2D to demonstrate, first, that all rats released along the long walls of the rectangular arena would end up at either corner b or d if they followed the IDF gradient and, second, that the same would happen with a proportion of rats released at the short walls. Rats were released with equal probability at these locations, and our conjecture was that on release they had to turn either along the wall or away from it. Inspection of Figure 2D suggests that release at the long walls thus accounts for 50% of rats ending up in corners b and d. If we assume the turning radius of an adult rat to be approximately 20 cm, rats released at the short walls will enter the catchment areas of corners b and d in about 10% of cases, bringing the predicted proportions of corners b and d entry to 60%. This is close to the proportion of first choices by rats in the first session of testing (Pearce et al., Figure 3, right panel), which we measured from their graph to be approximately 73% for both groups of rats combined, with the "inconsistent" group's choices counted as 100% minus the percent values given in the right panel of Figure 3 of Pearce et al. Simple inspection of IDF gradients thus does allow one to assess how an agent would move when released at any location within the arena and when "hunting" for a minimum of the IDF, which in this case does predict some aspects of rat search behavior.

To go beyond this qualitative assessment, we modeled this situation quantitatively. The optimal search model had parameters for inertia and random noise. These parameters were not estimated but were taken from the estimates in modeling Figure 3 of Graham et al. (2006; described in Experiment 2). The results of 100,000 simulated runs produced 60% correct choices, which is about the same estimate as that obtained with a qualitative analysis. The discrepancy between the performances of rats and the model stems largely from the poor performance of the simulated agent when released from the center of a short wall. Pearce et al. (2004, p. 140) reported that the rats headed for the nearest correct corner 60% of the time, whereas our simulated agent headed to the nearest corner only 37% of the time. When released from the center of a long wall, rats in Pearce et al. were about equally likely to head to either of the correct corners, whereas our simulated agent headed predominantly to the nearest correct corner from the release point. These discrepancies suggest that the rats had more information available than our reconstruction of views has identified, and that the simple model has likely missed something that rats do. To develop and test more detailed and realistic models, however, it is
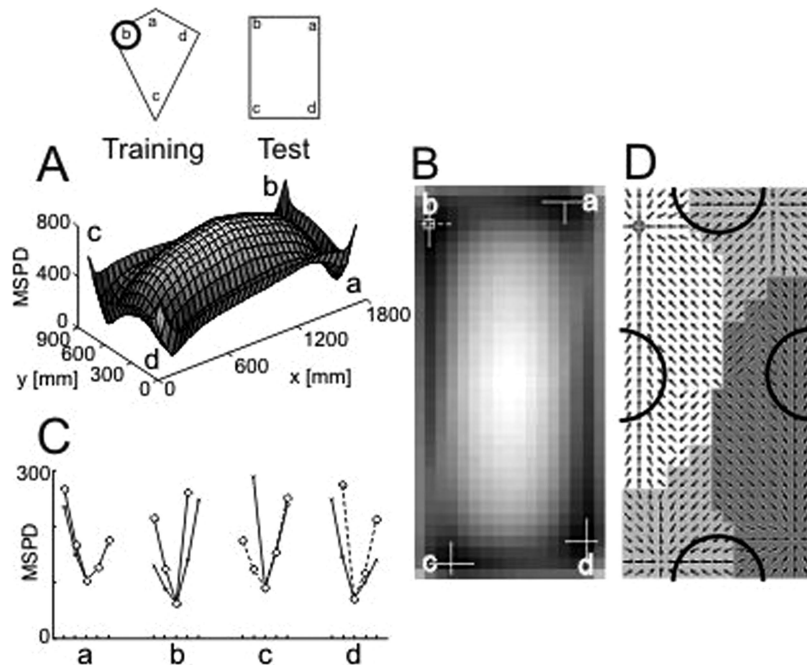
Figure 2. The image difference function (IDF) in the rectangular test arena after training in the kite-shaped pool at corner b. Conventions as in Figure 1. (A) Three-dimensional representation of the IDF. (B) Two-dimensional representation of the IDF. (C) Local transects through the local minima of the IDF. (D) Local slopes of the IDF and the catchment areas of local minima. Semicircles indicate start locations. Modeled after Pearce et al. (2004).

necessary to have detailed descriptions both of the actual visual scene and of how the rats move in initial orientation and in subsequent search.
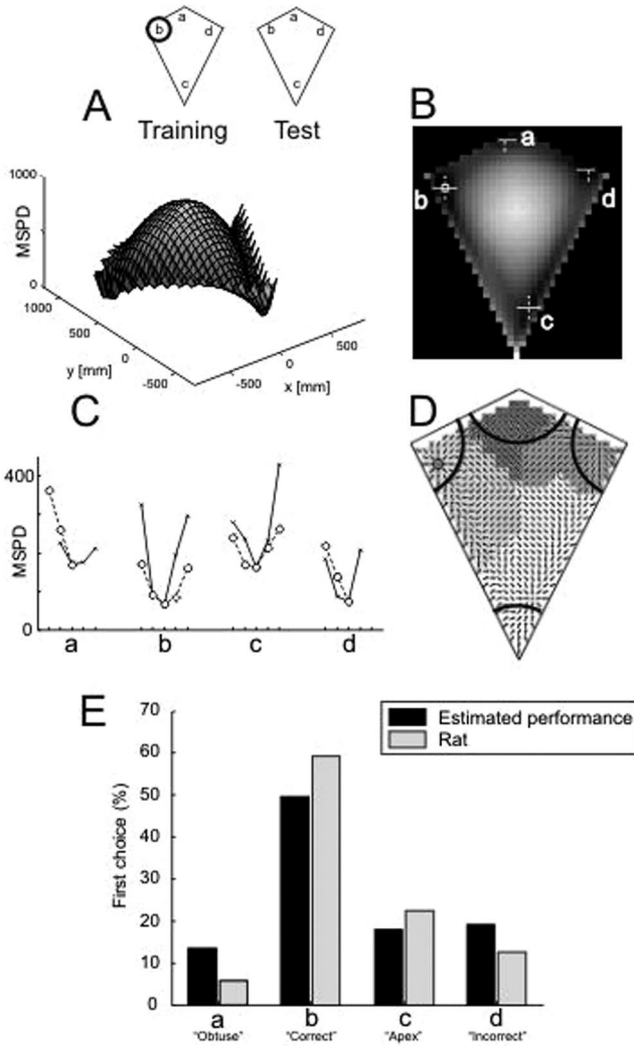
Pearce et al. (2004) showed that at the end of training in the kite-shaped pool, rats swam first to the training corner b about 60% of the time (their Figure 3, left panel). This is interesting because the arena at location b and d should look very similar. Simulating the IDF for this situation bears this out (Figures 3A–C): The *x*- and *y*-transects in Figure 3C show that the minimum IDF values at the target location b and at location d are about equally low. However, the catchment areas associated with these two locations tell a different story (Figure 3D) because the training corner has a much larger catchment area than its diagonal opposite. This goes some way toward explaining the performance of the rats with regards to corners b and d, but it does not account for the rats' ability to also avoid corner c, near the apex, which has a substantial catchment area (Figure 3D). Note, however, that the minimum associated with location c lies outside a 40-cm radius drawn around the apex (shown in black in Figure 3D), which was the criterion for corner entry used by Pearce et al. The criterion thus may underrepresent the visitation rate of the apex corner, leading to the reported low percentage. For example, assume that approximately half of those rats that reach the local minimum near the apex end up entering the apical arc first, and the rest end up in the next closest catchment area, that is, that of corner b. The first entry results based on the IDF catchment areas would then be similar to the experimental results (Figure 3E).

The preceding method of comparing simulation results with rat behavior, however, is somewhat arbitrary. We therefore

further investigated this last series of results of Pearce et al. (2004) by modeling in detail the performance of random and of view-driven search for this test situation. Model (based on 1,000 runs each) and rat performance are compared in Figure 4, with the model agent release points and corner entry areas shown in the inset on the top right. The first result to note is that search driven by pure noise leads the agent with equal probability of about 30% into corners a, b, and d, but only 10% of the time into the apex corner c (see black bars in Figure 4). This is very different from what rats do (white bars), which is to avoid the obtuse corner a and most frequently visit corner b. The second result to note is that both modes of view-based search, regardless of whether search is driven only by the IDF gradient (pure IDF) or with some added noise (optimal), mimic rat behavior much more closely than random search (dark and light gray bars in Figure 4). However, there remains a quantitative discrepancy because rats are clearly more attracted to corner b and less attracted to corner a than our modeling would predict. All we can surmise at this stage is that the rats may make use of some additional information that, for instance, could be nonuniform ceiling light distribution causing corner a to be more dissimilar to corner b than our virtual images suggest. It is, of course, possible that rats also use some other strategy in addition to image matching.

Experiment 2: Simulations of Graham et al. (2006)

In Experiment 1 of Graham et al. (2006), three groups of rats were trained in a kite-shaped arena and tested in the kite-shaped or
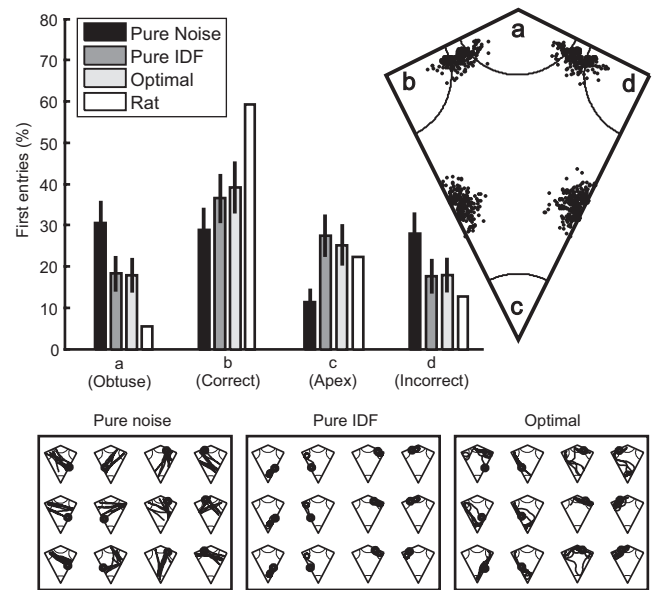
the target corner was always black; these rats were trained with only one view at the target corner. We simulated the information content of panoramic views in these conditions, including transfer tests in an all-black kite-shaped or a black and white square arena.

Method

The construction of virtual reality arenas was as before (see Experiment 1). The gray-level values of wall colors were picked from a set of real images, with only black walls forming the rectangular box and with only white walls forming the rectangular box. All other procedures were as in Experiment 1, with one exception: In simulations of experiments in which, during training, the submerged platform was located in two visually distinct environments (like in the shape-only and color-only conditions, see above), the following procedure was adopted. A reference image was taken at each of the two target locations. During sampling,


Figure 3. The image difference function (IDF) in the kite-shaped pool at the end of training to corner b in the same pool. (A–D) Conventions as in Figure 2. The arcs in D indicate the 40-cm radius used by Pearce et al. (2004) as criterion for entry to a corner. (E) Comparison of rat first-choice behavior (Pearce et al., 2004, Figure 3, left panel) with simulation results, based on the relative sizes of catchment areas shown in D. Modified from Pearce et al. Results were measured from their graph using graphics software. Because the local minimum near corner c falls outside of the 40-cm radius, half the arrivals there are assumed to end up first at corner b rather than corner c. Modeled after Pearce et al.

Figure 4. Comparison of performance model results with rat search behavior reported by Pearce et al. (2004) at the end of training in a white kite-shaped swimming pool. Shown is the performance for three model parameter settings: pure noise, pure image difference function (IDF), and optimal (see *Method* section, Experiment 1). For each parameter setting, the simulated agent is released 1,000 times with equal probability facing each of the four walls. Actual release locations are distributed as a Gaussian function with radial standard deviation of 10 cm (see inset top right), and release heading direction is distributed with the mean direction facing the wall and a Gaussian error distribution with standard deviation of 11.5°, which is independent of release position. The inset shows the corner demarcations defined by Pearce et al.; the dots represent 1,000 release sites with superimposed Gaussian noise. Note that because of the limited spatial resolution of the IDF, some release sites are deemed invalid even though they are within the geometric definition of the kite. Search performance was measured by determining the percentage of first entries into each corner of the enclosure (defined by a circular arc of radius 40 cm centered at the kite vertices). The panels below the bar graph show examples of simulated paths in the kite-shaped arena for the three model search modes. Each column represents three random examples of the same release site (with added Gaussian noise). A dot represents starting position.

a square arena. The target corner was a right-angled corner in training. For the shape-only group, the walls forming the target corner were white on some trials and black on other trials; the rats were thus trained with two different views at the target corner. In our simulation, we assumed that rats are able to keep at least two panoramic reference images simultaneously in memory. For the color-only group, the target corner might be either of the right-angled corners, but black walls always formed the target corner; these rats were also trained with two different views at the target corner. For the shape + color group, the target corner was always one of the right-angled corners, and the color of the walls forming

each reference image was independently compared with the current image. The two resulting MSPD values were subsequently summed. The final IDF was thus the sum of the two IDFs generated by the two distinct reference images; consequently, the MSPD values were higher for simulations in which there were two reference images. Thus, in these simulations, we assumed that the animals would not solve a conditional discrimination problem, in which they first discriminate which space they are in and then behave according to the requirements of that space. Rather, we assumed that the animals would not know which arena they are in and, hence, would use both the reference images on all trials.

Modeling performance. For performance modeling (see Figure 10, which will be discussed in the *Results and Discussion* section), simulated agents were released at the center of the kite-shaped and the square test arenas, much like experimental rats in the study by Graham et al. (2006), with a joint Gaussian probability density function (radial standard deviation of 20 cm)

and with random orientation. On the basis of the performance criteria of Graham et al., the percentage of time that each simulated agent spent in each of four quadrants was found for each of 1,000 simulated paths for each experimental condition. The quadrants were defined by the arena's long axis of symmetry and a perpendicular line bisecting it. Consequently, a large quadrant of the kite-shaped arena constituted exactly 11/32 or approximately 34.4% of the total area (not 38.2%, as reported in Graham et al.).

Choosing parameters. The model contains two parameters to be fitted, a noise level ($0 \leq \alpha \leq 1$), and a heading filter size (an angle σ). To avoid "overfitting," we chose the two parameters on the basis of the 12 data points in Figure 3 of Graham et al. (2006). Mean squared errors, (model – rat)^2, on the 12 data points were obtained from 1,000 runs at each combination of parameter values, at a resolution of 0.1 for each parameter. The best fitting parameters were $\sigma = 0.4$ radians and $\alpha = 0.1$. Using absolute differences between model and data on the difference between "correct" and "incorrect" corners resulted in the same best parameters. These
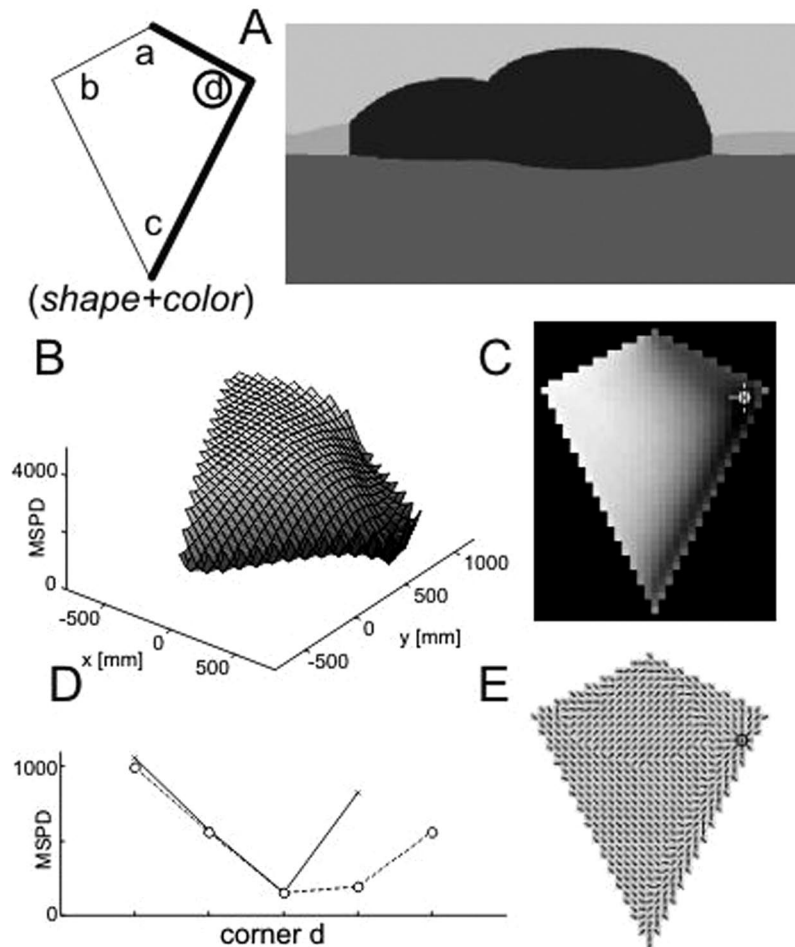
Figure 5. Experimental setup of the shape + color condition of Graham et al. (2006). Rats were trained in a kite-shaped pool with two black and two white walls. During training, the target corner was specified by having black walls and by always being located at the same corner d (see schematic diagram in A). (A) Simulated view of the training corner. (B) Three-dimensional representation of the image difference function (IDF) for this training location. (C) Two-dimensional representation of the same IDF. (D) Local transects through the only minimum of the IDF at the target corner d. (E) The catchment area of each of the local minima.

parameters were used for all subsequent modeling, including that already presented in Experiment 1.

Results and Discussion

Experiment 1 of Graham et al. (2006) had three training conditions for rats in a kite-shaped pool with black and white walls. In one condition (shape + color), one long wall and one short wall were black, and the other pair was white. Walls of each color formed a right-angled corner, and the target corner was black (see schematic diagram in Figure 5A). The simulated reference image at the training location d shows a large black expanse with the white walls in the background (Figure 5A). The IDF for this situation shows one clear minimum at the training corner, with no other local minima present (Figures 5B–D). The catchment area for this location covers the whole test space so that an agent hunting for the minimum of image differences would find the correct corner from anywhere in the pool (Figure 5E). And indeed, rats reach 100% performance in this task after some training, as

documented by Figure 2 of Graham et al. Our performance model yielded 981 first entries into the correct corner out of 1,000 trials (98.1%).

In a second condition (color-only), rats could use wall color cues but not shape cues to locate the target. The target might be at either of the right-angled corners b or d of the kite-shaped space, but the walls about the target are always black (see insets of Figure 6A). Following a view-based strategy, an agent would acquire two reference images at these alternative locations (Figure 6A). We simulated the situation at the end of training for two black walls on the right of the arena (inset of Figure 6B). The IDF shows three local minima (Figure 6C); two are close to the target location for this set up (d, d′ in Figure 6C) and have a lower minimum value (Figure 6D) than the one that lies along the long wall toward the apex (c′ in Figure 6C). The catchment areas in Figure 6E show that the local minimum along the long wall at location c′ is very close to the catchment area of location d, so that small searching movements would lead to entry into the catchment area associated with
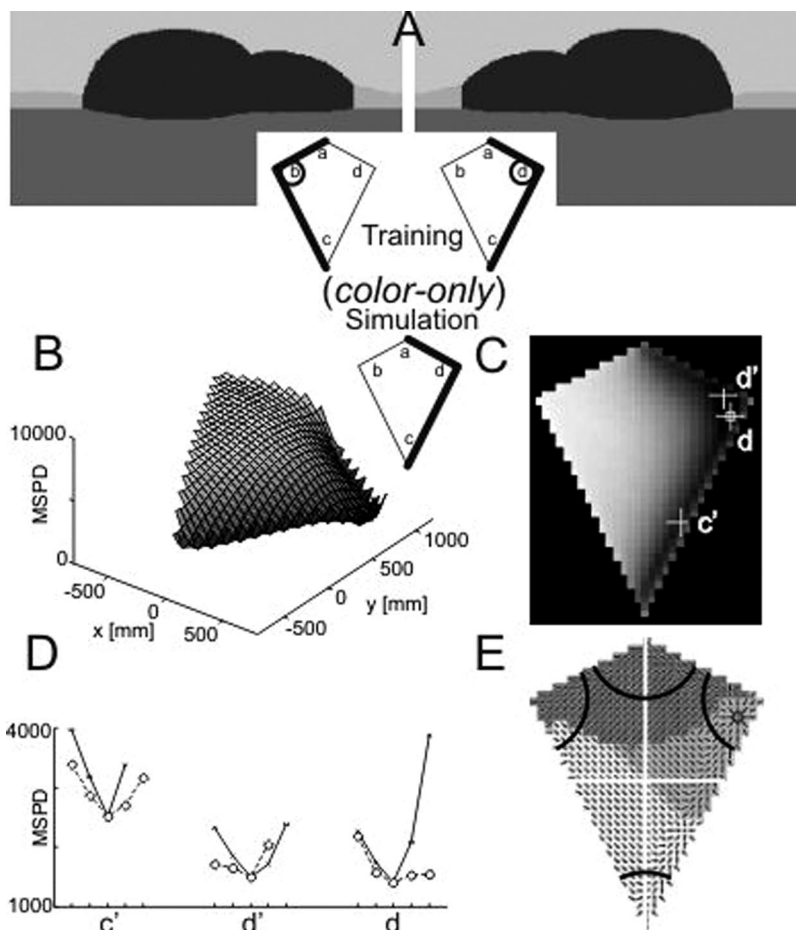

Figure 6. Color-only condition of Graham et al. (2006). Rats were trained to rely exclusively on wall color cues (see inset). (A) Simulated views at the two training corners. (B) Three-dimensional representation of the image difference function (IDF) for the test situation with black walls on the right side of the kite-shaped pool (see inset). (C) Two-dimensional representation of the same IDF. (D) Local transects through the minima of the IDF. (E) Local slopes of the IDF and catchment areas based on the two training views. The white lines show the division of the space into the quadrants defined by Graham et al.

this location. Unfortunately, we cannot test this conjecture against experimental data because Graham et al. did not report entries to the apex but only compared entries into corners b and d. We thus would predict on the basis of our simulation of catchment areas that performance should be 100% (always d before b), the value found for rats (Figure 2 of Graham et al.). Our performance model yielded 978 first entries into the correct corner out of 1,000 trials (97.8%).

The third condition of training (shape-only) forced the rats to rely on geometry alone and to ignore conflicting wall color information. The target was always at one corner (the right-angled corner on the right of the upright kite-shaped arena), but it might have two white walls flanking it (left panel of Figure 7A) or two black walls flanking it (right panel of Figure 7A). Intuition would suggest that navigation by image matching would prove difficult under such circumstances, and our simulations showed just that.

For the situation with two black walls on the right, the combined IDF does not have pronounced minima (Figures 7B and 7C), but a multitude of shallow ones (Figure 7D), all of which have similar, high remaining image difference values. The only noticeable feature is that the strongest gradient of image differences leads to the

minimum at the target location d′ (Figure 7D). The map of catchment areas in Figure 7E shows that one minimum, a′, lies between the obtuse corner a and the corner diagonally opposite the target at b. Another minimum, b′, lies near b but outside of the 40-cm radius.

We used our performance model to simulate this experiment. Our simulated rats visited the target corner first on 738 out of 1,000 trials (73.8%). In Session 20, shape-only rats in Graham et al. (2006) visited the target corner first on 70% of the trials (as measured from their Figure 2 using graphics software). We conclude that the training data from Experiment 1 of Graham et al. can be explained by postulating that rats navigate in these experimental spaces with the aid of remembered views.

Graham et al. (2006) also conducted unrewarded tests on each of the trained groups in their Experiment 1. In one test, rats were placed in a kite-shaped pool that was all black. This was meant to test whether rats are using geometric information. Color was made less informative in that black walls now formed all corners. We determined the IDF for this test, with the goal at the right-angled corner on the right of the upright kite shape. The results for the shape + color group trained with both shape and color cues (i.e.,
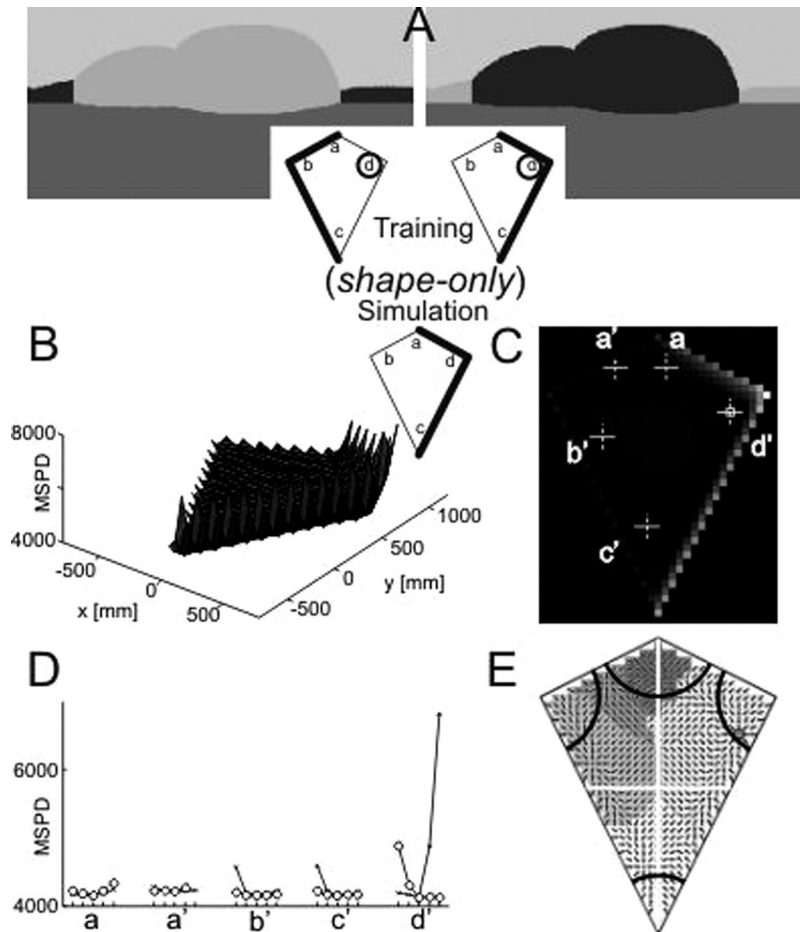

Figure 7. Shape-only condition of Graham et al. (2006). Rats were trained to ignore wall color cues and to rely on shape cues only (see inset). (A–E) Conventions as in Figures 5 and 6. The arcs in E indicate the 40-cm radius used by Graham et al. to define entry into a corner.

with two black walls, two white walls, and the goal at the corner with two black walls) are shown in Figure 8-I. The reference image is shown in Figure 5A. The IDF for this situation has clear local minima near each corner, with the one at the apex displaced along the long wall on the left (Figures 8-IA and B). The lowest image differences are located at the goal corner d and the diagonally opposite corner b (Figure 8-IC). The catchment area of the goal corner d is larger than the one of the diagonally opposite corner b (Figure 8-ID). If we assume for a qualitative assessment that the size of catchment areas translates into time spent in each quadrant, this simulation result does explain what rats seem to do: Graham et al. recorded the quadrant visited by rats and found that rats preferred to visit the quadrant containing the target corner d (Figure 3, left panel, of Graham et al.).

The IDF of the second treatment (color-only), in which rats were trained with consistent wall color at the goal but inconsistent location (Figure 8-II) showed unsurprisingly that the appearance of the goal corner and the opposite corner are very similar. The reference images for this test are shown in Figure 6A, with their IDF in Figures 8-IIA and B. The IDF has six local minima, two of which lie near the goal corner d (d, d′ in Figures 8-IIB and C) and two close to the opposite corner b (b, b′ in Figures 8-IIB and C). The local minimum near the apex (c) has a much higher value than all the others (Figure 8-IIC). The catchment areas are similar for the goal corner d (catchments for d and d′ added together) and the

opposite corner b (catchments for b and b′ added together; Figure 8-IID) because each matches one of the reference images. On the basis of memorized images, there is no way to discriminate between the two corners, and rats are also unable to do so (Graham et al., 2006, Figure 3, left panel).

Finally, we determined IDF properties in an all-black kite by the shape-only group, which had a goal corner that was sometimes black and sometimes white (Figure 8-III). The two reference images are shown in Figure 7A. The IDF (Figures 8-IIIA and B) again reveal four local minima near each of the corners, similar to Figures 8-IA and B, but all showing high remaining image differences and negligible differences between them (Figure 8-IIIC). This leads to a slightly larger catchment for the top right quadrant compared with the top left one (Figure 8-IIID). Rats visit the two quadrants at about the same rate (Graham et al., 2006, Figure 3, left panel).

Graham et al. (2006, Experiment 1) also tested the rats trained in the kite-shaped pool in a square arena. This arena had two adjacent black walls and two adjacent white walls. Corners could thus be discriminated by wall color but not by shape information. Rats of the shape + color training group predominantly visited the black quadrant of the square test box, as did the color-only group, whereas the shape-only group did not show such a preference (Graham et al., Figure 3, right panel).
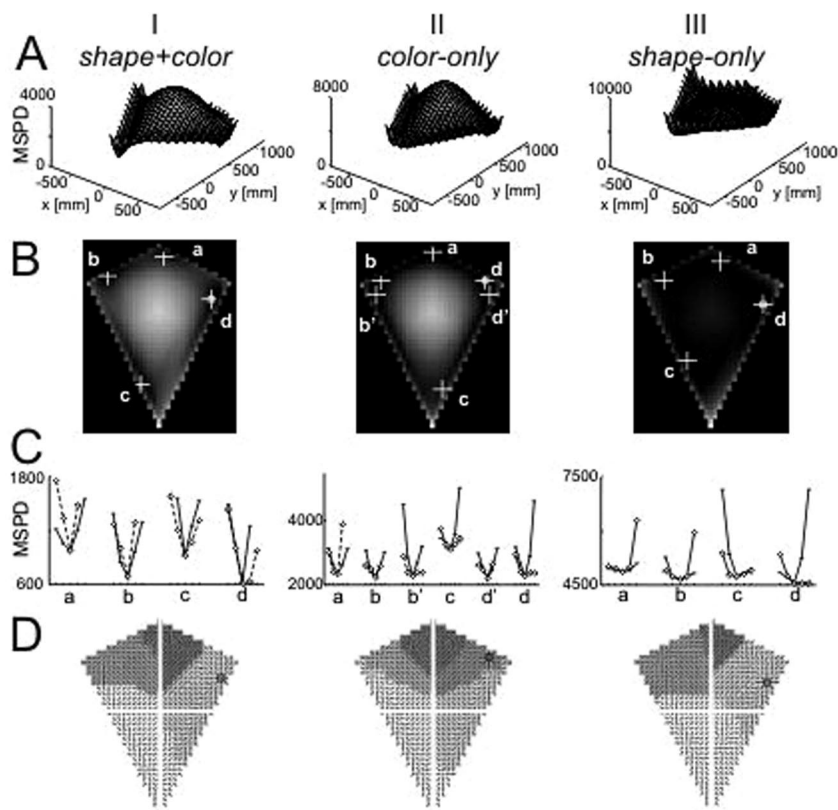

Figure 8. Rats trained in shape + color, color-only, and shape-only conditions were tested in an all-black kite-shaped pool (Graham et al., 2006). The virtual reality image difference functions, transects, and catchment areas for these three conditions are shown in rows A–D for the three conditions. Conventions as in Figure 2. The white lines show the division of the space into the quadrants defined by Graham et al. (2006).

The properties of the IDF for this test suggest that the search behavior of rats may be driven by remembered views (see Figure 9). The IDF of the test for the shape + color group (Figure 9-I; reference image Figure 5A) reveals a pronounced local minimum at the goal corner d, which has a large catchment area practically filling the goal quadrant. On the basis of catchment area size, one would thus predict a clear preference of rats for corner d. The test situation for the second color-only group (Figure 9-II; reference image Figure 6A) produces an IDF with two dominant, closely spaced minima in corner d whose catchment areas practically cover the whole test space, leading to a clear, predicted preference for that corner. In contrast, the IDF for the shape-only group (reference image Figure 7A) in the square box shows multiple shallow local minima with large remaining image differences (Figure 9-III). The catchment areas of the "correct" (all black) corner d at the top right and the "incorrect" (all white) corner at b, b' are of about equal size and larger than the other two catchment areas. On the basis of catchment area size, we would thus expect that rats would visits the two corners b and d about equally often.

We simulated search performance for the situations shown in Figures 8 and 9 and compared them with the performance of rats shown in Figure 3 of Graham et al. (2006). For both the experiments in the kite-shaped and the square arenas, the results of performance modeling are surprisingly similar to the search behavior of rats (see Figure 10). The bar graph shows percent time spent in the "correct" quadrant in black and percent time spent in the "incorrect" quadrant in white for rats having received shape + color, color-only, and shape-only training. The performance modeling was done with optimal model parameter settings (optimal search; see the *Method* section, Experiment 1). Although the performance model predicts the search pattern of rats across these very different treatments extremely well qualitatively, we again found that rats trained in the color + shape condition do quantitatively better in both the kite-shaped and the square arenas test by entering the "correct" quadrant more often than the model would predict. This again suggests that the experimental environment, but not our virtual environment, may have offered additional cues to corner location. To what extent this conjecture also explains the two remaining deviations between model and rat search (color-only, kite-shaped arena and shape-only, square arena) is at this stage hard to decide. The current performance model is no doubt overly simplistic, but it does serve two functions. First, it shows an example of how local IDF values, if available, may be used to search for a goal. Second, it shows that using the IDF gives the simulated agent a better chance of finding the "correct" location than random searching (modeled by the pure noise condition) in
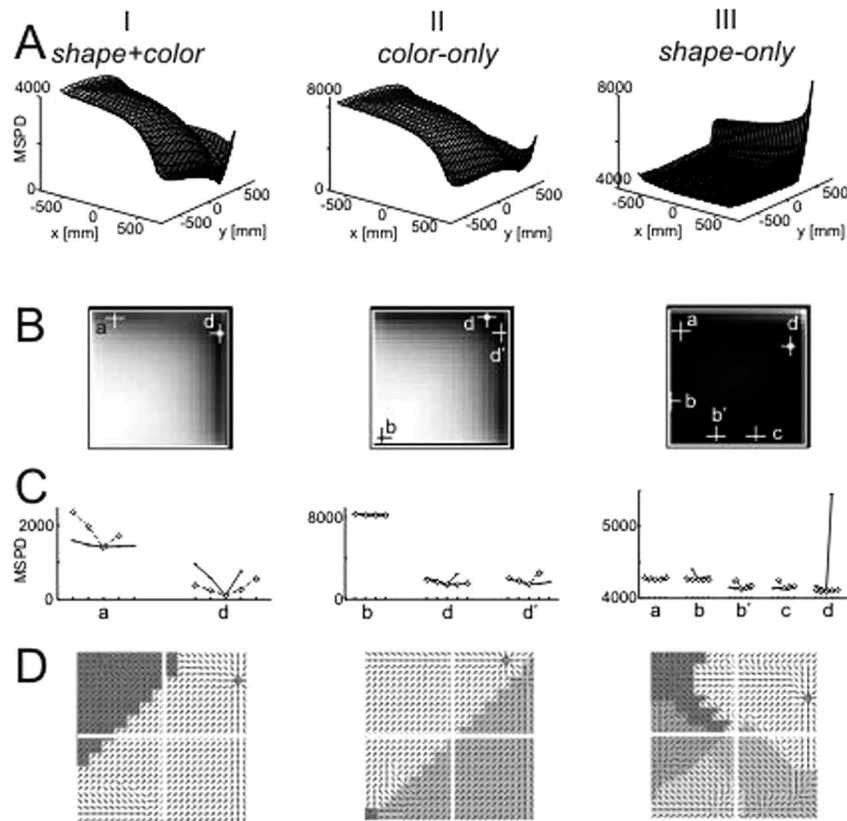

Figure 9. Rats trained in shape + color, color-only, and shape-only conditions in a kite-shaped pool were tested in a square pool with two black and two white walls (Graham et al., 2006). The virtual reality image difference functions, transects, and catchment areas for these three conditions are shown in rows A–D. Conventions as in Figure 2. The white lines show the division of the space into the quadrants defined by Graham et al. (2006).
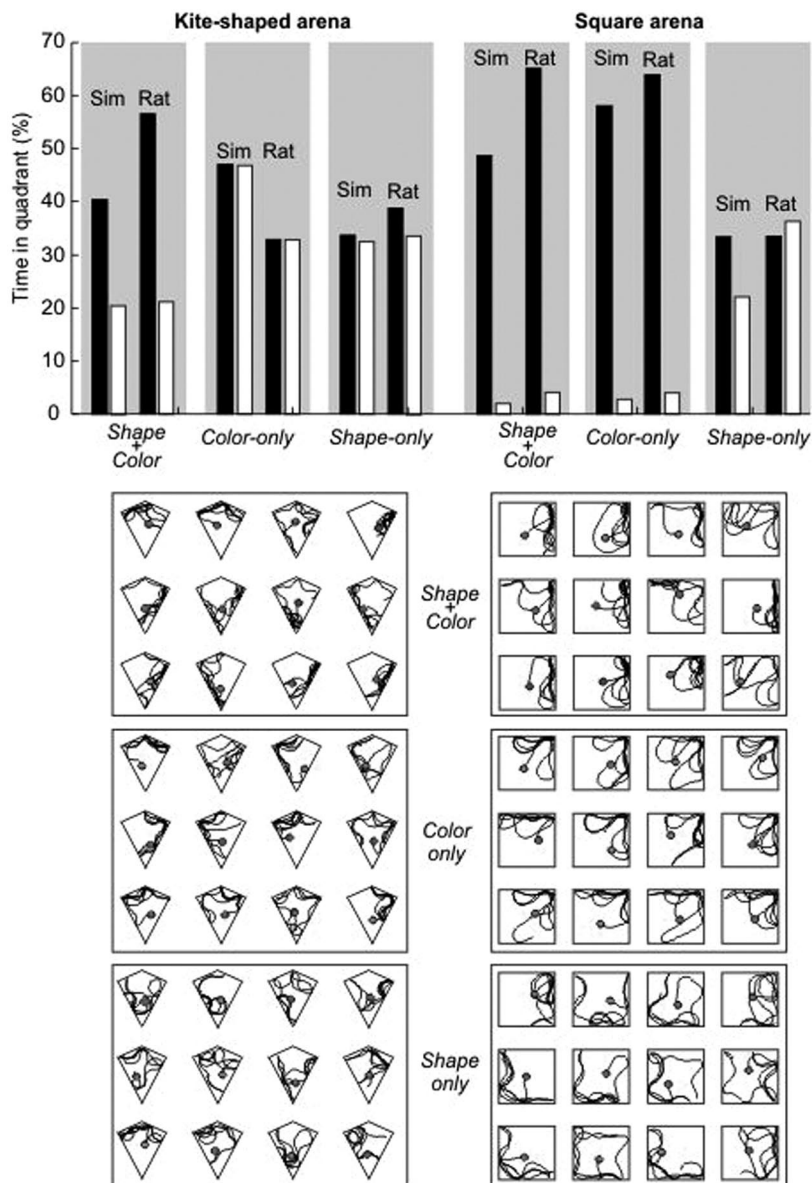
Figure 10. Comparison between rat choice behavior (data from Graham et al., 2006, Figure 3) and model search performance for training in an all-black kite-shaped arena and test in a kite-shaped (top left) and a square arena (top right). Shown are the results of minimizing the absolute difference between "correct" minus "incorrect" percentages for model and rat performance. Optimal model parameters for noise ($\alpha = 0.1$) and heading filter width ($\sigma = 0.4$ rad) were used in a modeling condition called *optimal search* (see *Method* sections, Experiments 1 and 2). In each condition, the left pair of bars represents the performance modeling results (Sim), and the right pair represents rat data (Rat). The black bars represent percent time spent in the "correct" quadrant, and the white bars represent percent time in the "incorrect" quadrant. The panels below the bar graph show examples of simulated paths with optimal model parameters for the kite-shaped (left) and the square arena tests (right). Each column represents three random examples of the same release site (with added Gaussian noise). A dot represents starting position.

both training and test arenas and approximates more closely the behavior of rats.

General Discussion

We explored whether the search behavior of rats in swimming pools of different shapes (Graham et al., 2006; Pearce et al., 2004) can be explained by a relatively simple view-based homing mechanism. Our results indicate that the search behavior of rats can in principle be explained by assuming that the animals take a snapshot at the goal corner of, for instance, a rectangular training space and that they move in such a way during tests in a kite-shaped arena as to minimize the difference between the memorized image

and the views experienced in the test arena. Such an image-matching procedure does not require extraction of edges, objects, or spatial geometry. By determining the catchment areas of such panoramic snapshots and by comparing a random search strategy with search driven by image differences, we have shown that the latter mimics the search performance of rats closely across very different training and test regimes.

The results of our simulations are a long way from showing that rats rely exclusively on a visual matching strategy to navigate in simple experimental spaces, such as the swimming pools typically used. Rats in all likelihood may use a number of different search strategies, including random or systematic search (Bartumeus, da Luz, Viswanthan, & Catalan, 2005), but they also may have more cues at their disposal than simple image comparison to expedite escape from the water. Our results, however, suggest a most parsimonious and easily testable explanation for a wide range of experimental data.

The results shown in Figure 3 of Graham et al. (2006) and our Figure 10 are especially noteworthy. As Graham et al. wrote in their general discussion, the pattern poses difficulties for extant learning theories. One possibility is that rats extract some geometric properties of the kite-shaped experimental arena and use their representation to navigate in the test spaces. Such cues might be local properties such as corner angles and relevant wall lengths (Pearce et al., 2004; Tommasi & Polli, 2004). An example for solving the kite problem in Graham et al. might be to pick a right-angled corner with a long wall on the right. Identifying a right angle and discriminating the lengths of walls are necessary for this solution. The geometric properties might also be based on parameters extracted from the shape of the arena (Cheng, 2005; Cheng & Gallistel, 2005). An example would be to pick out the longitudinal axis and go to the corner to the right of this axis at the wide end of the kite. The main problem faced by these geometry-based hypotheses is to explain why the wall colors should have such an impact. The results of Test 1 of Experiment 1 by Graham et al., (2006), reproduced in our top panel of Figure 10, indicate that the shape-only condition cannot be solved in the all-black kite-shaped arena, neither by rats nor by a simulated view-based homing algorithm. This test rendered color cues far less reliable, whereas shape-based geometric cues were unambiguously valid. According to theories that claim the extraction and use of geometric properties (Cheng, 2005; Cheng & Gallistel, 2005; Cheng & Newcombe, 2005; Gallistel, 1990; Pearce et al., 2004; Tommasi & Polli, 2004), shaped-based geometric cues, local or global, ought to especially stand out and be relied on. The data of Graham et al. (2006) contradict such expectations. What we suggest here is that these geometric cues are contained implicitly in panoramic images and that their salience relative to internal features will determine to what extent an animal's search is driven by one or the other cue (see Stürzl et al., 2007).

Another common strategy in learning theory, on the basis of Rescorla and Wagner (1972), would conceive of geometry and color as independent cues competing for associative strength in learning (Chamizo, 2003; Mackintosh, 2002). One upshot of this kind of cue competition is that with an additional redundant cue (color) in learning, a cue of interest (shape) would not be learned as well. The all-black kite test of Graham et al. (2006) was designed to eliminate color cues and to force animals to attend to shape cues only. Any cue competition theory would predict that the shape + color group, with a redundant color cue in competition with shape cues, ought to learn shape less well than the shape-only group, which was forced to rely exclusively on shape cues during training. The results, in the top panel of Figure 10, are just the opposite, with the shape + color group doing better. Both rats and our simulations show this pattern. The additional redundant color cue aided shape learning rather than hindering it, so the argument went in the general discussion of Graham et al.

Graham et al. (2006) conceived of the pattern of their data as indicating *potentiation* and provided a discussion of various ways in which this might take place. Descriptively, this means that having an additional redundant cue aids the learning of a cue in question. A well-documented example is the learning of an odor aversion, which is aided rather than hindered by having a distinctive taste associated with the odor (Durlach & Rescorla, 1980; Rusiniak, Palmerino, Rice, Forthman, & Garcia, 1982). In the potentiation explanation, shape and color of the target corner are extracted and then associated with one another, so that one cue can help the other.

Potentiation, or facilitation of the learning of geometric cues by the learning of featural cues, can also arise in an adaptation of the Rescorla–Wagner model for learning geometry and features (Miller & Shettleworth, 2007). Featural and geometric cues, as well as cues provided by food bowls, are conceived of as separate kinds of cues, analogous to different conditioned stimuli in classical conditioning. According to Miller and Shettleworth (2007), the critical difference between the operant spatial-search task and classical conditioning is that in the former, the rat (or any other subject) gets to choose the stimuli it is exposed to on a trial. When the rat learns a unique featural cue at the correct corner, both the featural cue and the geometric cues at the corner gain associative strength. In this way, learning the correct featural cue can often increase the learning (associative strength) of geometric cues.

An explanation in terms of view-based navigation seems to be more parsimonious, in that it does away with the need to separately learn, store, and associate shape and color cues. Nor is there a need for a dichotomous mixture of cue potentiation and cue competition to explain different experiments as suggested by Graham et al. (2006). Rather, in the particular simulations that we conducted, the overall view is simply recorded as a panorama of pixel values, which are subsequently used for image comparison, without feature extraction or object recognition and a fortiori without extraction of geometric properties. As the results demonstrate, such a simple strategy goes a long way toward explaining the pattern of results found by Pearce et al. (2004), Graham et al., and others (see Stürzl et al., 2007). Indeed, we have also found many results of Tommasi and Polli (2004) to be well explained by this strategy (data not shown).

Miller and Shettleworth's (2007) model rescues the notion of cue competition by showing that within a Rescorla–Wagner framework adapted for the learning of geometry and features, cues may sometimes facilitate the learning of other cues and at other times interfere or compete with the learning of other cues. A view-based matching strategy also rescues the notion of cue competition in landmark use. That is because only one kind of cue, a panoramic image, is at play, thus taking cue competition out of the arena, metaphorically speaking. Our results should not be taken as evidence against cue competition in landmark use. Rather, if view-based matching is the strategy used, cue competition is not tested because there is only one kind of cue. In other situations of spatial

learning, much evidence supports the claim that individual landmarks are subject to cue competition in rats (Chamizo, 2003) and other animals (Cheng & Spetch, 1998, 2001). Whether a hypothesis of view-based matching can predict this cue competition remains to be seen. It is possible, for example, that some cases of blocking in landmark use might reflect a failure or slowness to update a reference image.

It also remains to be seen how much of the "geometry" literature in rats (Cheng & Newcombe, 2005) can be accounted for by a view-based homing strategy. But what it can account for so far is worth noting. A view-based strategy can explain the original results obtained by Cheng (1986), which were most favorable to an account in terms of the use of geometric properties (Stürzl et al., 2007). And now we have shown that a view-based strategy accounts just as well for results that are highly troublesome for an explanation that requires the extraction of geometric cues. We suspect that view-based matching can also account for the ease with which rats learn the position of a place relative to an apparatus, even when it conflicts with the Earth-based position of the place within a room (a situation created by translating the apparatus in the room; Hamilton, Akers, Weisand, & Sutherland, 2007; Skinner et al., 2003). Hamilton et al. (2007), for example, trained rats to find a hidden platform located in one quadrant of a round pool situated in a room with many extrapool cues. When the pool was translated within the room, so that the target position relative to the pool (using distal cues as a frame for directions only) and target position relative to the room conflicted, the rats preferred to search at the target position relative to the pool rather than the target position relative to the room. The authors concluded that directional information is obtained from the distal cues beyond the pool, whereas distance information is obtained from the pool itself. From the perspective of view-based matching, however, the target position relative to the pool is probably the position with the least mismatch to the reference image learned during training. At this position, the walls of the pool and the expanse of water, which form a large part of the panoramic view, would line up correctly; the mismatch comes from a shift in the distal cues. At the target position relative to room, on the other hand, vast expanses of wall and water would not match what was seen at the target location during training. In general, the closer the landmarks are to the target location, the more their influence on the IDF (Stürzl & Zeil, 2007). View-based modeling should be done to confirm these intuitions about the results of Hamilton et al.

It should be emphasized that this work represents only a first step toward a full mechanistic and quantitative model of animal search behavior in experimental arenas or in natural space. We have purposefully kept all modeling extremely basic, incorporating just a few core ideas. Among many things, we have left out all considerations of learning a view as well as relearning or updating views. The modeling basically assumes a fully detailed reference image, and only animals at asymptotic performance were modeled. With all these limitations, it is surprising that the outcome of our simulations seems to be consistent with such a wide range of results both qualitatively and quantitatively. It is possible that elaborations of our simple view-based navigation model, perhaps incorporating biomechanical constraints, view-independent directional information, or neurobiological data, may widen the explanatory scope of this general approach.

For quite some time, experiments on insect homing have suggested that view-based navigation forms a key part of insect navigational strategies (e.g., Cartwright & Collett, 1983, 1987, 2002; Cheng, 2006; Collett & Collett, 2002; Collett & Zeil, 1998; Wehner, 2003). It is increasingly being recognized that image differences offer robust cues to location in natural environments (Stürzl & Zeil, 2007; Zeil et al., 2003) and therefore can also be usefully employed in robotic navigation (Franz, Schölkopf, Mallot, & Bülthoff, 1998; Vardy & Möller, 2005). Our results suggest that image matching can parsimoniously explain how rats search for places of interest in confined spaces. It may also form a part of the navigational repertoire of some other vertebrates whose spatial abilities have been well documented (Cheng, Spetch, Kelly, & Bingman, 2006; Etienne & Jeffery, 2004; Healy, 1998, 2004; Rodriguez et al., 2002).

References

Bartumeus, F., da Luz, M. G. E., Viswanthan, G. M., & Catalan, J. (2005). Animal search strategies: A quantitative random-walk analysis. *Ecology, 86,* 3078–3087.

Cartwright, B. A., & Collett, T. S. (1983). Landmark learning in bees. *Journal of Comparative Physiology A, 151,* 521–543.

Cartwright, B. A., & Collett, T. S. (1987). Landmark maps for honeybees. *Biological Cybernetics, 57,* 85–93.

Cartwright, B. A., & Collett, T. S. (2002). How honeybees use landmarks to guide their return to a food source. *Nature, 295,* 560–564.

Chamizo, V. D. (2003). Acquisition of knowledge about spatial location: Assessing the generality of the mechanism of learning. *Quarterly Journal of Experimental Psychology, 56*(B), 102–113.

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,* 5–21.

Cheng, K. (2006). Arthropod navigation: Ants, bees, crabs, spiders finding their way. In E. A. Wasserman & T. R. Zentall (Eds.), *Comparative cognition: Experimental explorations of animal intelligence* (pp. 189–209). Oxford, England: Oxford University Press.

Cheng, K., & Gallistel, C. R. (2005). Shape parameters explain data from spatial transformations: Comment on Pearce et al. (2004) and Tommasi and Polli (2004). *Journal of Experimental Psychology: Animal Behavior Processes, 31,* 254–259.

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

Cheng, K., & Spetch, M. L. (1998). Mechanisms of landmark use in mammals and birds. In S. Healy (Ed.), *Spatial representation in animals* (pp. 1–17). Oxford, England: Oxford University Press.

Cheng, K., & Spetch, M. L. (2001). Blocking in landmark-based search in honeybees. *Animal Learning & Behavior, 29,* 1–9.

Cheng, K., Spetch, M. L., Kelly, D. M., & Bingman, V. P. (2006). Small-scale spatial cognition in pigeons. *Behavioural Processes, 72,* 115–127.

Collett, T. S., & Collett, M. (2002). Memory use in insect visual navigation. *Nature Reviews Neuroscience, 3,* 542–552.

Collett, T. S., & Zeil, J. (1998). Places and landmarks: An arthropod perspective. In S. Healy (Ed.), *Spatial representations in animals* (pp. 18–53). Oxford, England: Oxford University Press.

Durlach, P. J., & Rescorla, R. A. (1980). Potentiation rather than overshadowing in flavor-aversion learning: An analysis in terms of within-compound associations. *Journal of Experimental Psychology: Animal Behavior Processes, 6,* 175–187.

Etienne, A. S., & Jeffery, K. J. (2004). Path integration in mammals. *Hippocampus, 14,* 180–192.

Franz, M. O., Schölkopf, B., Mallot, H. A., & Bülthoff, H. H. (1998). Where did I take that snapshot? Scene-based homing by image matching. *Biological Cybernetics, 79,* 191–202.

Gallistel, C. R. (1990). *The organization of learning.* Cambridge, MA: 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,* 44–59.

Hamilton, D. A., Akers, K. G., Weisand, M. P., & Sutherland, R. J. (2007). How do room and apparatus cues control navigation in the Morris water task? Evidence for distinct contributions to a movement vector. *Journal of Experimental Psychology: Animal Behavior Processes, 33,* 100–114.

Healy, S. (Ed.). (1998). *Spatial representation in animals.* Oxford, England: Oxford University Press.

Healy, S. H. (2004). Spatial learning and memory in birds. *Brain, Behavior and Evolution, 63,* 211–220.

Jeffery, K. J. (Ed.). (2003). *The neurobiology of spatial behaviour.* Oxford, England: Oxford University Press.

Mackintosh, N. J. (2002). Do not ask whether they have a cognitive map but how they find their way about. *Psicologia, 23,* 165–185.

McGregor, A., Jones, P. M., Good, M. A., & Pearce, J. M. (2006). Further evidence that rats rely on local rather than global spatial information to locate a hidden goal: Reply to Cheng and Gallistel (2005). *Journal of Experimental Psychology: Animal Behavior Processes, 32,* 314–321.

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

Morris, R. G. M. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation, 12,* 239–260.

Newcombe, N. S. (2002). The nativist–empiricist controversy in the context of recent research on spatial and quantitative development. *Psychological Science, 13,* 395–401.

Olton, D. S., & Samuelson, R. J. (1976). Remembrance of places passed: Spatial memory in rats. *Journal of Experimental Psychology: Animal Behavior Processes, 2,* 97–116.

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,* 135–147.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcment and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64–99). New York: Appleton-Century-Crofts.

Roberts, W. A. (1984). Some issues in animal spatial memory. In H. L. Roitblat, T. G. Bever, & H. S. Terrace (Eds.), *Animal cognition* (pp. 425–443). Hillsdale, NJ: Erlbaum.

Rodriguez, F., Lopez, J. C., Vargas, J. P., Gomez, Y., Broglio, C., & Salas, C. (2002). Conservation of spatial memory function in the pallial fore-brain of reptiles and ray-finned fishes. *Journal of Neuroscience, 22,* 2894–2903.

Rusiniak, K. W., Palmerino, C. C., Rice, A. G., Forthman, D. L., & Garcia, J. (1982). Flavor–illness aversions: Potentiation of odor by taste with toxin but not shock in rats. *Journal of Comparative and Physiological Psychology, 96,* 527–539.

Skinner, D. M., Etchegary, C. M., Ekert-Maret, E. C., Baker, C. J., Harley, C. W., Evans, J. H., et al. (2003). An analysis of response, direction, and place learning in an open field and T maze. *Journal of Experimental Psychology: Animal Behavior Processes, 29,* 3–13.

Stürzl, W., Cheung, A., Cheng, K., & Zeil, J. (2007). The information content of panoramic images: I. Rotational errors and the similarity of views in rectangular experimental arenas. *Journal of Experimental Psychology*: *Animal Behavior Processes.*

Stürzl, W., & Zeil, J. (2007). Depth, contrast and view-based homing in outdoor scenes. *Biological Cybernetics, 96,* 519–531.

Tommasi, L., & Polli, C. (2004). Representation of two geometric features of the environment in the domestic chick (*Gallus gallus*). *Animal Cognition, 7,* 53–59.

Vardy, A., & Möller, R. (2005). Biologically plausible visual homing methods based on optical flow techniques. *Connection Science, 17,* 47–89.

Wang, R. F., & Spelke, E. S. (2002). Human spatial representation: Insights from animals. *Trends in Cognitive Sciences, 6,* 376–382.

Wang, R. F., & Spelke, E. S. (2003). Comparative approaches to human navigation. In K. J. Jeffery (Ed.), *The neurobiology of spatial behaviour* (pp. 119–143). Oxford, England: Oxford University Press.

Wehner, R. (2003). Desert ant navigation: How miniature brains solve complex tasks. *Journal of Comparative Physiology A, 189,* 579–588.

Zeil, J., Hofmann, M. I., & Chahl, J. S. (2003). Catchment areas of panoramic snapshots in outdoor scenes. *Journal of the Optical Society of America, 20,* 450–469.