nd enlarging upon the discussion of
ally in the sections on planning and
raditionally thought to be uniquely
es and ask whether they depend on
nd birds, for example, understand

8

Getting Around: Spatial Cognition

Limpets are small mollusks that live on coastal rocks where they are exposed to the air at low tide. As it grows, each limpet erodes a scar on the rock that matches the irregular outline of its shell. By clinging tightly to this spot during low tide, the limpet can protect itself from dehydration (and from predators, as anyone who has tried to pry one loose can testify), but to find food it must forage over the rock while the water is high (Cook et al. 1969).

Many mobile animals face the same problem as the limpet: food and other resources are separated from places of refuge, and the animal has to be able to travel between them without getting lost. There is a premium on making this trip efficiently rather than wandering at random until the goal is found, which in the limpet's case might be too late to prevent drying out. The limpet's problem is a miniature one in space and time compared to the orientation problems solved by other species (Figure 8.1), but they all have certain features in common. Each individual or group of individuals is locating its own home, hoards, or other resources. Therefore, they need some sort of acquired representation of the goal's location or how to get to it. Some animals, like the limpet, create such a representation in the external world in the form of a chemical trail (Chelazzi 1992). Under some circumstances, animals can find their way by directly approaching cues emanating from their goals or learning sequences of responses. However, we will be most concerned with how animals acquire and use information that is inherently spatial, that is, information about distances and directions. Mathematically, this is vector or shape information. And although long-distance navigation may involve amazing feats of perception, learning, and memory (see Box 8.1), we will be almost entirely concerned with travels of a meter or less to at most a few kilometers.

Because acquiring and acting on spatial information appears to have different computational requirements from learning to predict temporal sequences of events, we might expect to find adaptively specialized, domain-specific mechanisms of spatial learning and/or performance, different from those for associative learning (F. Dyer 1998; Gallistel 2003).This issue can be addressed in terms of the three fundamental questions about learning from Chapter 4—the conditions for learning, the contents of learning, and its effects on behavior. Spatial performance rules can be thought of as servomechanisms. That is, they generate behavior that reduces the discrepancy between the animal's current position and a remembered target position (Cheng 2000). Although mechanisms for

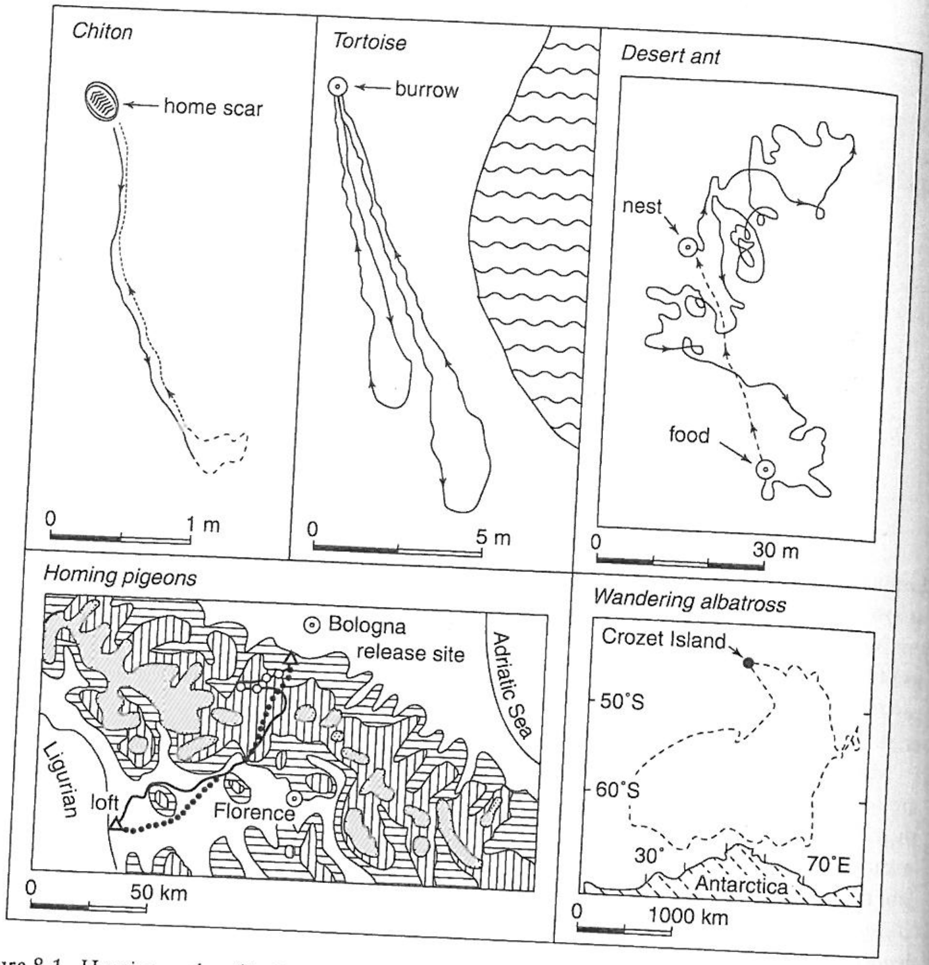

Figure 8.1. Homing paths of individuals from five species, illustrating the wide range of spatial scales over which journeys out from a central place and back may take place. After Papi (1992b). Tortoise from Chelazzi and Francisci (1979); albatross from Jouventin and Weimerskirch (1990). Redrawn with permission.

Box 8.1 Long-Distance Migration

The astonishing ability of animals from all taxa to find their ways over hundreds or thousands of kilometers is a subject in itself (see Alerstam 2006; Holland, Wikelski, and Wilcove 2006). It has been given a huge boost in recent years by sophisticated satellite tracking systems for recording not only the position but the activities, temperatures, and so forth, of migrating animals. The sensory and neural mechanisms required can also be studied in some of the same species (Frost and Mouritsen 2006). Notwithstanding their vastly different scales, however, long-distance and short-distance travels are largely analyzed with the same basic conceptual framework (Bingman and Cheng 2005). Distance, direction, and position information are important however far one is going, and the degree to which it is maplike is an issue whatever its scale.

Probably the longest-standing subjects in studies of the mechanisms for long-distance migration are birds. Among the many species of small birds that migrate at night, even captive hand-reared individuals exhibit nocturnal activity, so-called migratory restlessness, at the time of year when they would normally migrate (see Gwinner 1996). In indoor cages at night they tend to hop toward the compass direction in which their conspecifics are flying at that time of year. Manipulating the early experience of such birds has revealed a kind of interaction between predisposition and experience that might be called *calibration*. Calibrating a physical measuring instrument means comparing its readings to those of an independent standard and adjusting it so its readings match the standard's. An

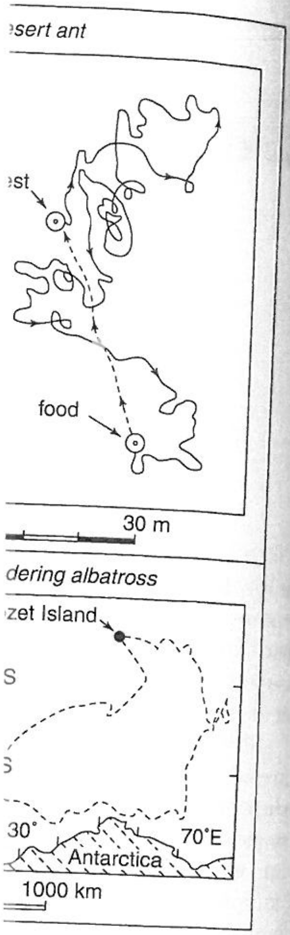

g the wide range of spatial
e place. After Papi (1992b).
n and Weimerskirch (1990).

r hundreds or thousands of
d Wilcove 2006). It has been
tems for recording not only
; animals. The sensory and
ecies (Frost and Mouritsen
istance and short-distance
ork (Bingman and Cheng
vever far one is going, and

or long-distance migration
even captive hand-reared
the time of year when they
ey tend to hop toward the
ar. Manipulating the early
lisposition and experience
nent means comparing its
s match the standard's. An

---

electronic thermometer might be calibrated against a mercury thermometer, for example. Analogously, one orientation mechanism may be changed by experience so that its outputs more closely match those of a second, independent, mechanism. The primary examples involve calibrating celestial cues against magnetic information.

For example, the primary directional cues for nocturnal migrants are the Earth's magnetic field and, on clear nights, the stars, but the pattern of stars varies with geographic location, time of night, and season, and it changes over geologic time. Insight into how birds nevertheless use the stars to tell direction comes from classic experiments by Emlen (1970) with indigo buntings (*Passerina cyanea*). He raised three groups of birds indoors out of sight of the sky, but late in their first summer two of those groups were exposed to the "night sky" in a planetarium. For one, the stars rotated normally, around the North star, whereas for the other the center of rotation was the bright star Betelgeuse. When all the birds then spent autumn nights in the planetarium under stationary star patterns typical for the time of year, the birds with no experience of the sky were not well oriented, but those exposed to the normal sky oriented Southward, indicating that they had somehow learned to use the stationary star patterns during earlier exposure to the normal night sky. The third group treated Betelgeuse as the North star, flying "south" with respect to it, indicating that the star or star pattern near the center of rotation of the night sky is used to give direction. Magnetic information interacts with this information during normal development (see Able and Bingman 1987; Able and Able 1990; Weindler, Wiltschko, and Wiltschko 1996).

Some species change direction in midjourney, following routes that take them around inhospitable places like the Alps and the Sahara. Young birds raised in captivity show evidence of population-specific genetic programs that specify the duration of migratory restlessness and its direction with respect to the magnetic field (Helbig 1994, 1996). Figure B8.1 shows an example in which two European populations of a single species, the blackcap (*Sylvia atricapilla*), migrate in different directions, and one changes course part way while the other does not. Such inborn tendencies to head in a certain compass direction at a certain season are likely important for the many species in which animals migrating for the first time are not accompanied by experienced adults. This likely includes sea turtles and at least one insect, the Monarch butterfly (Holland, Wikelski, and Wilcove 2006). At the same time, the success of programs for reintroducing migratory bird species to their ancestral flyways by training them to follow ultralight aircraft (www.operationmigration.org) indicates that some species learn details of their migratory routes.

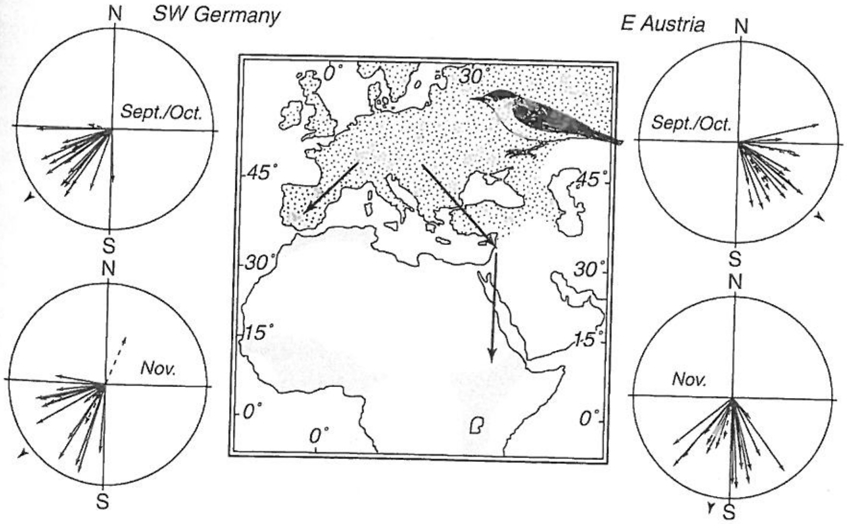

Figure B8.1. Breeding (dotted) and wintering areas (grey) of two populations of European blackcaps with their major migration routes. Arrows inside the circles represent orientation of hand raised birds from the two populations tested in funnel cages early and late during the period when they would normally be migrating. Each vector is the mean orientation of a single bird; the longer the arrow the stronger the directional tendency. Redrawn from Helbig (1994) with permission.

reading out where to go from information about where one is are far from simple (Biegler 2006), the effects of spatial learning on behavior are generally taken for granted: the animal reaches the goal in the presence of the appropriate cues. Much more attention has been devoted to the content of spatial learning. In practical terms, this means discovering what features of the goal control behavior. The most controversial question about the content of spatial learning is "Do animals have cognitive maps?" That is, is spatial orientation in complex environments controlled by an overall representation of distances and directions that allows the animal to select an efficient route when displaced to a new location? This question turns out to be difficult to answer, for two reasons. First, although a map is a powerful metaphor for spatial knowledge, different investigators may mean different things by *cognitive map*. Second, before we can consider whether any animal might have a cognitive map in any sense, we need to consider all the simpler mechanisms animals can use to find their ways to goals (Section 8.1). and how they may be combined (Section 8.2). Section 8.3 discusses how animals acquire spatial knowledge, especially whether any processes different from associative learning are involved. Then, in Section 8.4, we will assess the evidence for cognitive maps.

8.1 Mechanisms for spatial orientation

8.1.1 Dead reckoning

A foraging desert ant (*Cataglyphis fortis*) wanders here and there, taking a long and tortuous path in its search for food, but as soon as it finds a prey item it heads straight back to its nest over a hundred meters away (see Figure 8.1). These ants return to the vicinity of the nest using *dead reckoning,* an internal sense of the direction and distance of the nest from their current position. That they know both distance and direction can be shown by catching an ant in a matchbox just before it starts its homeward journey and releasing it several hundred meters away. It does not head for the nest but takes a path parallel to that which it would have taken from the point of capture. For instance, if the nest was originally to its south, the ant still heads south even if the nest is now to the east (Wehner 1992, 2003). Moreover, when it has gone about the right distance, the ant begins to circle around as if looking for the nest in the place where it should be (Figure 8.2). This behavior shows that the ant must be performing *path integration* on the outward journey. That is, it behaves as if continuously integrating (in the mathematical sense) information about its changes in distance and direction to compute the vector that links it to the nest.

In fact, ants use an approximation in which each direction taken, as perceived by its solar compass (Box 8.2) is weighted by the distance for which it is maintained (Muller and Wehner 1988). The orientation of the straight path reveals the ant's representation of the homeward direction, and the point at which it begins to circle around reveals its representation of the distance from start to nest. Once the ant arrives near where the nest should be, it continues to perform the same implicit computations. Although taking a roughly spiral path, it continually returns to the point where it began searching, as if keeping track of its position with respect to the most likely nest position. This localized search seems to be programmed to overcome the inherent errors of path integration in that the further an ant has traveled from the nest, the wider its spiraling loops when it returns to the nest's vicinity (Wehner and Srinivasan 1981; Merkle, Knaden, and Wehner 2006). This behavior increases the chances that the nest is found, which is vital because the hot sand surface can be lethal to ants that do not escape underground quickly enough.

far from simple (Biegler
taken for granted: the
luch more attention has
, this means discovering
sial question about the
ps?" That is, is spatial
resentation of distances
hen displaced to a new
reasons. First, although
nvestigators may mean
er whether any animal
ne simpler mechanisms
they may be combined
knowledge, especially
lved. Then, in Section

re, taking a long and
item it heads straight
se ants return to the
f the direction and
v both distance and
t before it starts its
It does not head for
n from the point of
int still heads south
r, when it has gone
g for the nest in the
at the ant must be
behaves as if con-
bout its changes in
est.
en, as perceived by
h it is maintained
reveals the ant's
it begins to circle
est. Once the ant
the same implicit
lly returns to the
vith respect to the
med to overcome
traveled from the
nity (Wehner and
vior increases the
face can be lethal

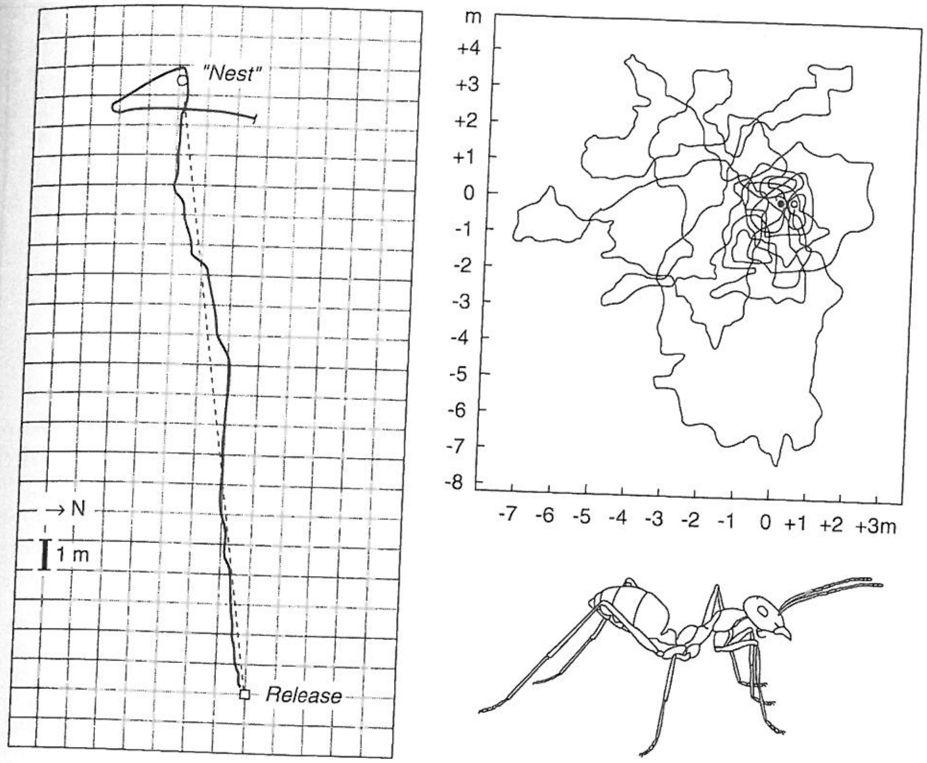

Figure 8.2. Homing in desert ants (*Cataglyphis albicans*). Left: Path of an individual that has just found food and is displaced to unknown territory. The open circle shows where the nest would have been relative to the release site if the ant had been in its home territory. Right: The spiraling path taken by the ant once it arrives near where the nest should be, recorded over one hour. Data redrawn from Wehner and Srinivasan (1981) with permission; ant from Wehner (1992) with permission.

Box 8.2 The Sun Compass

The sun is useless as a landmark because it moves continuously relative to the Earth, but many diurnal animals use it for directional information, that is, they have a *sun compass*. For example, the desert ants in Section 8.1.1 use both the sun and patterns of polarized light it creates in the sky for directional information when computing their paths home from food (Wehner and Müller 2006). If an ant is trained to a food source on a featureless patch of desert at one time of day and then kept in the dark for a few hours, it heads roughly homeward when released even though its direction relative to the sun's position is different from what it was during training. That they are still relying on the sun rather than some subtle landmarks is shown by the fact that ants prevented from seeing the sun in this experiment head off in random directions (Wehner and Lanfranconi 1981).

Reading direction from the sun regardless of the time of day requires both a stored representation of how the sun moves across the sky at the current location and season (an *ephemeris function*) and an internal circadian clock (Chapter 9). The sun's position overhead is converted to a compass direction (i.e., direction relative to North) by computing the sun's *azimuth*. This means taking the imaginary arc connecting the sun with the closest point on the horizon and measuring the angle on the surface of the earth between that point and North (Figure B8.2a). This kind of computation is implied by statements like "The sun is in the South" at noon in the Northern hemisphere. But although the sun *is* in the South at noon, because the sun's elevation at a given time of day changes with the time of year, the sun's azimuth changes at different rates at different times of year and at different times of day (Figure B8.2a). Thus to use the sun for directional information, animals must acquire some

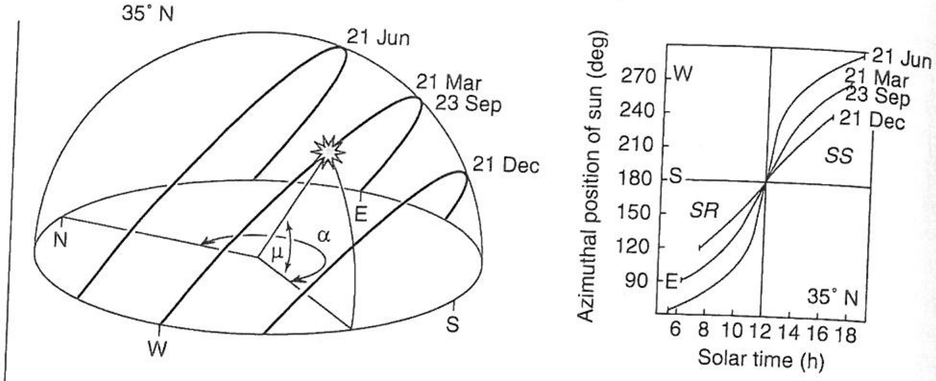

Figure B8.2a. How the apparent path of the sun across the sky (arcs) varies with time of year at a particular latitude, 35° North. Angle α on the surface of the earth is the sun's azimuth; μ is the sun's elevation. Ephemeris functions (right) give the sun's azimuth as a function of time of day and time of year. After Wehner (1992) with permission.

representation of the local ephemeris function and continually update it. This process has been studied in honeybees by restricting the experience of newly hatched foragers when they first leave the hive (F. Dyer and Dickinson 1994). The results indicate that, in a kind of process general to many kinds of learning, bees begin life with a crude default ephemeris function, a best guess about the conditions they are likely to meet, and experience fine-tunes it (F. Dyer and Dickinson 1996).

To show definitively that an animal is using a sun compass it is necessary to shift its internal clock and test whether orientation shifts accordingly. (As discussed in Chapter 9, shifting the clock means keeping the animal under an altered day-night light cycle for several days.) The logic of clock shift experiments is depicted in Figure B8.2b with a hypothetical example using bees. Homing pigeons have also been tested extensively in such experiments (Papi and Wallraff 1992). Of course in laboratory studies of small-scale spatial learning, animals cannot use a sun compass because the

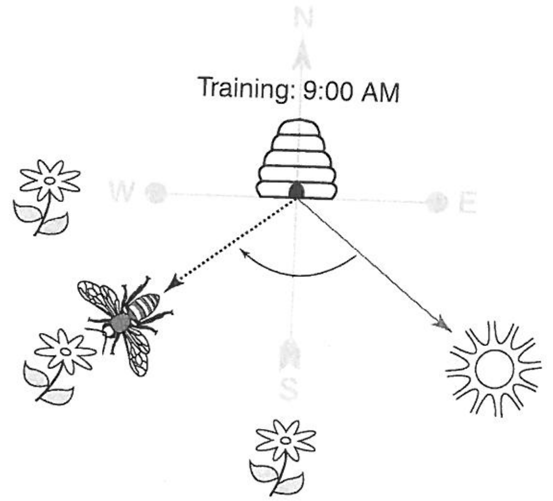

Figure B8.2b. The logic of clock shift experiments, showing how to tell which way a clock-shifted animal will head. In this example, a bee trained to find food in the position indicated has its clock shifted back 3 hours by turning the lights on 3 hours later in the morning. At 9 AM it experiences the time as 6 AM. When it flies from the hive it will maintain the same angle to the sun as when heading to the goal at 6 am before clock shifting.

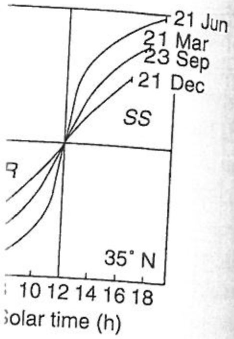

ies with time of year at
sun's azimuth; μ is the
unction of time of day

This process has been
ragers when they first
, in a kind of process
ephemeris function, a
fine-tunes it (F. Dyer

shift its internal clock
ifting the clock means
he logic of clock shift
bees. Homing pigeons
1992). Of course in
mpass because the

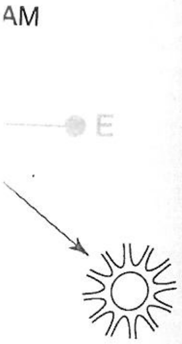

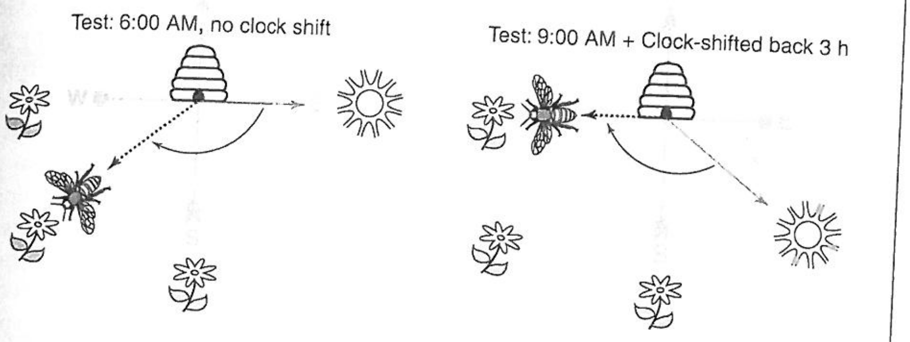

Figure B8.2b. (Continued)

sun is not visible. However, some birds have proved to use the sun compass in learning simple spatial discriminations outdoors under sunny skies. These include homing pigeons (Bingman and Jones 1994; Chappell and Guilford 1995), scrub jays (W. Wiltschko and Balda 1989), and black-capped chickadees (Sherry and Duff 1996). For example, when scrub jays were clock-shifted by 6 hours between storing and retrieving seeds in an outdoor arena, the birds relied on their sun compass in spite of the fact that distant landmarks were visible outside the arena. However, the relative importance of the sun compass vs. other spatial information will vary with species and circumstances as suggested by the discussion of pigeons' homing in the main text.

Dead reckoning is one of the most basic and ubiquitous ways in which animals keep track of their location with respect to a known position. (*Dead reckoning* is a navigators' term; it is generally used interchangeably with *path integration.*) It has been studied most in insects such as ants, bees, and spiders (Wehner and Srinivasan 2003). Indeed, although possible examples of dead reckoning in humans were noted by Darwin (1873), its role in spatial learning by rats and other small mammals was almost completely overlooked before Mittelstaedt and Mittelstaedt (1980) described it in gerbils (*Meriones unguiculatus*). In the situation they studied, mother gerbils and their pups had a nest at the edge of a large circular arena. If the pups were taken from the nest and placed in a cup somewhere in the arena, the mother soon began to search for them. When she found the pups, she picked one up in her mouth and ran almost straight back to the nest, even in total darkness and even if her outward path had zigzags and detours. If the nest was moved by rotating the edge of the arena while the mother was at the stationary cup, she returned to the starting point of her journey like the desert ant, ignoring any cues emanating from the nest in its new location. In contrast, if the cup was rotated briskly while the mother gerbil was in it, she compensated for the rotation and headed straight back to the nest as before. But if the cup was rotated slowly or slowly moved sideways, the gerbil did not compensate and was misoriented. The effect of rotation speed reflects that fact that in mammals information about changes in angular orientation is processed by the vestibular system, which senses accelerations and decelerations above a certain threshold (McNaughton, Knierim, and Wilson 1995; Wallace et al. 2002).

More extensive studies like these have been done by Etienne and her colleagues with golden hamsters (*Mesocricetus auratus*) hoarding food from the center of an arena back to their nest and increasingly with rats (Etienne and Jeffery 2004). Geese carried in a cart up to a kilometer or so from their home also appear to home by dead reckoning (Saint Paul 1982). They obtain information about displacement from the patterns of visual flow. If they cannot see out of the cart for parts of the outward journey, they act as if discounting this part of the trip. This intriguing little study has apparently never been followed up, and few further observations relevant to path integration in birds have been reported. Pigeons show little evidence of relying on visual flow for position information in a laboratory task (Sutton and Shettleworth 2005). In contrast, a great deal is known about how ants and bees compute distances and directions of travel from visual and other cues (Boxes 8.2 and 8.3). Reliance on nonvisual, vestibular, cues for direction is especially appropriate for nocturnal species like hamsters and rats. However, although the sensory inputs are very different in mammals and insects, the implicit computations on them are similar. For example, when forced to take an outward journey consisting of two segments connected at a given angle, ants, spiders, bees, and several species of mammals make similar angular errors when heading home (Etienne and Jeffery 2004).

Box 8.3. Odometers of Honeybees and Desert Ants

We see in the main text that honeybees and desert ants behave as if having an *odometer,* a mechanism for measuring distance traveled. But bees generally fly whereas ants walk, and the odometers of the two species use correspondingly different information. For flying honeybees, distance is measured by *optic flow,* the angular motion of images past the eyes. Evidence comes from experiments such as the one diagrammed in Figure B8.3a, in which bees flew down a tunnel decorated with vertical black and white stripes to find sugar water (Srinivasan et al. 1996). With the food at a fixed location, bees learn where to expect it as evidenced by their circling around over the usual place of food in unrewarded tests. When image motion was eliminated by replacing the vertical stripes by horizontal ones for the tests, the bees searched equally at all distances. When the tunnel was wider or narrower than usual, the bees searched at a greater or lesser distances respectively (Figure B8.3a). To understand why the effect of the tunnel's width, that is, the distance of images from the eyes, means that angular image motion is important, think of how nearby objects cross your visual field faster than those farther away when you are in a moving car. Changing the density of the pattern inside the tunnel also changes the rate of image motion, and accordingly, in natural landscapes the bees' subjective estimates of distance as revealed in their dances (Section 14.2.1) is greater when they have flown over a richly patterned landscape than when they have flown the same distance over water (Tautz et al. 2004).

Desert ants walk across rather featureless terrain. Accordingly they estimate distance using about the only cue available, the number of steps they have taken. In the most direct demonstration that the ant's odometer is in fact a pedometer, ants that had walked along a straight channel from the nest to food were captured before starting home and fitted with stilts made of pig bristle or made to walk on stumps by painlessly removing the last segment of their legs (Wittlinger, Wehner, and Wolf 2007). Then they were released in a long parallel test channel and—as in the tests with bees—the point at which they began circling around searching for the nest was recorded. The altered ants walked in a remarkably normal way. As a result those with stilts went too far, and those with stumps not far enough (Test 1 in Figure B8.3b). In contrast, ants that had stilts or stumps throughout a whole round trip estimated the nest location accurately (Test 2 in Figure B8.3b).

Normally ants compute a straight homeward path by path integration over a winding outward journey as in Figure 8.2. What if part of the journey is over hilly terrain? Remarkably, the ant's pedometer compensates for hills, perhaps relying on gravity sensors in the joints (Grah, Wehner, and Ronacher 2005). Ants that either left the nest over a hilly channel and were transferred to a flat

ne and her colleagues
from the center of an
d Jeffery 2004). Geese
pear to home by dead
isplacement from the
parts of the outward
iguing little study has
ions relevant to path
vidence of relying on
ton and Shettleworth
es compute distances
and 8.3). Reliance on
for nocturnal species
are very different in
imilar. For example,
nents connected at a
make similar angular

ving an *odometer,* a
s ants walk, and the
or flying honeybees,
eyes. Evidence comes
s flew down a tunnel
et al. 1996). With the
ling around over the
ted by replacing the
all distances. When
r or lesser distances
, that is, the distance
f how nearby objects
g car. Changing the
and accordingly, in
eir dances (Section
hen they have flown

istance using about
nonstration that the
nel from the nest to
or made to walk on
r, and Wolf 2007).
bees—the point at
ed ants walked in a
ith stumps not far
out a whole round

winding outward
arkably, the ant's
ts (Grah, Wehner,
ransferred to a flat

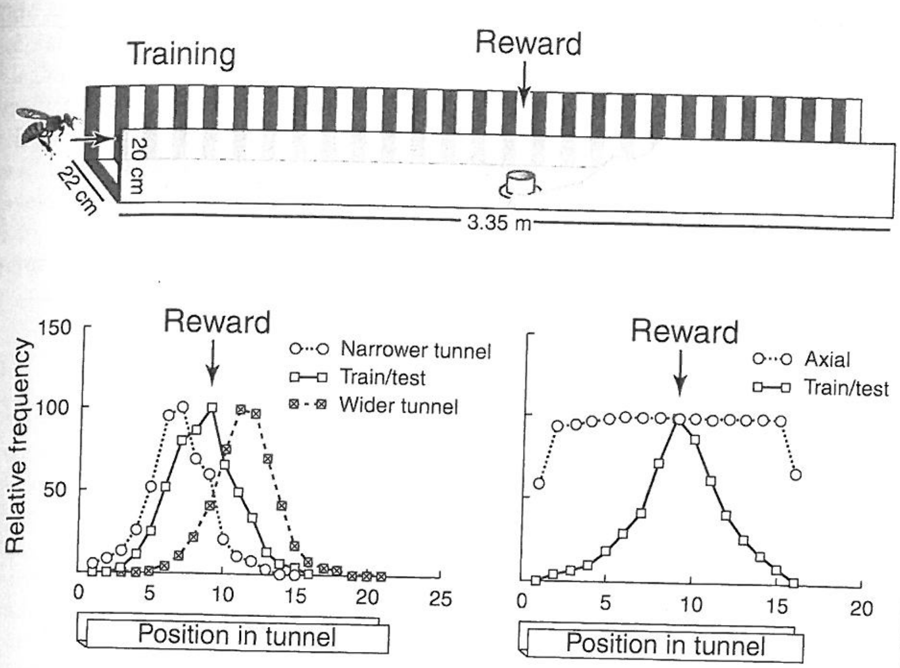

Figure B8.3a. Setup and results of experiment testing influence of visual flow on distance estimation in honeybees. The data are proportions of searches, normalized to 100 at the peak place of searching. The measure of position in the tunnel is number of vertical stripes. All bees were trained to the same position, the one used for trials marked "train/test" but then tested with wider and narrower tunnels or axial stripes in the same tunnel (right panel). Adapted from Srinivasan et al. (1996) with permission.

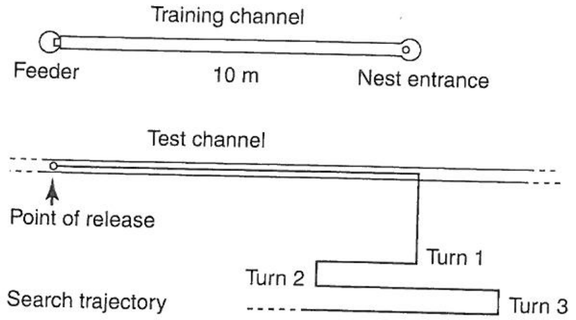

Figure B8.3b. Data from test of odometry in ants on stilts and stumps that was otherwise analogous to the study with bees in B8.3a. Adapted from Wittlinger, Wehner, and Wolf (2006) with permission.

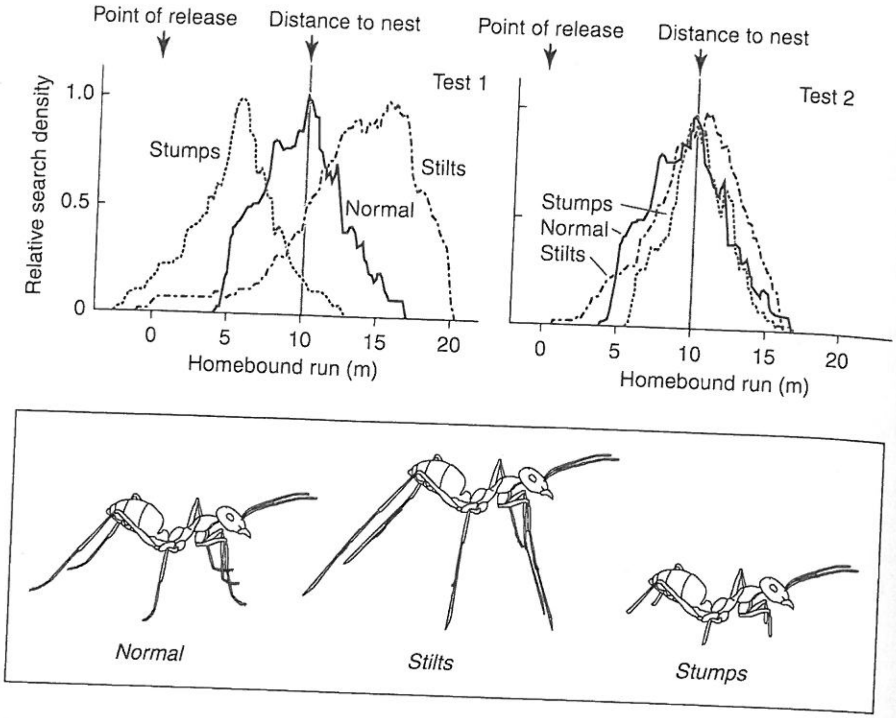

Figure B8.3b. (Continued).

channel to home or the reverse searched for the nest at the correct distance over the ground (Wohlgemuth, Ronacher, and Wehner 2001). And when ants that had traveled around a bend and over a steep "hill" to find food were released on open ground, they headed in the correct direction to find the nest and searched for it at the correct distance from the release point (Grah, Wehner, and Ronacher 2005). That is, they behaved like ants that had traveled to the same feeder over flat ground, not ants that had walked the same number of steps.

The similarities between Figures B8.3a and B8.3b imply that bees and ants compute distances using essentially the same implicit countinglike process but on qualitatively different inputs. We know very little about whether and how any mammals, for example nocturnal rodents, sense distance traveled as such. Most laboratory studies of path integration in rats or hamsters test primarily its directional component: in a confined space, animals can choose which way to head but have little choice in how far to go.

Dead reckoning is a mechanism for *egocentric* spatial localization, that is, the animal is localizing things in the environment with respect to itself. *Allocentric* (or *geocentric*) mechanisms locate the animal with respect to some external frame of reference such as landmarks or environmental geometry. We have already seen one of the major disadvantages of egocentric mechanisms: if the animal is slowly "blown off course," as by the experimenter moving it, path integration does not necessarily compensate. It also accumulates error. For instance, the more the hamsters have been turned around or have turned themselves around while collecting food from the center of the hoarding arena, the less accurately they return to the nest (Etienne, Maurer, and Saucy 1988).

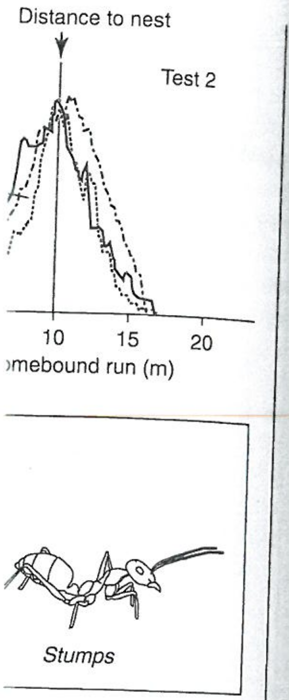

distance over the ground
d traveled around a bend
hey headed in the correct
e from the release point
that had traveled to the
mber of steps.
d ants compute distances
ively different inputs. We
nocturnal rodents, sense
in rats or hamsters test
oose which way to head

ocalization, that is, the
t to itself. *Allocentric*
ect to some external
l geometry. We have
ic mechanisms: if the
enter moving it, path
cumulates error. For
und or have turned
f the hoarding arena,
er, and Saucy 1988).

This makes dead reckoning most useful for comparatively brief round-trip excursions, as does the way it seems to be reset at the start of each new journey (Biegler 2000). A major advantage of dead reckoning is its availability from the first trip into a new part of the environment, before there has been time to learn reliable external cues. This makes it a potential basis for learning other cues. Dead reckoning is not only a one-trial affair, though. When hamsters repeatedly traveled in the dark on a circuitous path to a pile of food, they could still find it when forced to make a novel detour (Etienne et al. 1998), evidently using the vector computations of the dead reckoning system (Figure 8.3).

8.1.2 Beacons

In Mittelstaedt and Mittelstaedt's (1980) experiments, we might have expected odors or sounds from the nest itself to act as a *beacon* for the mother gerbil returning with a wandering pup. Beacons are sometimes referred to in the psychological literature as *proximal cues,* that is, cues close to the goal, as distinct from *distal cues,* the landmarks to be discussed in the next section. (*Local* vs. *global* cues is much the same distinction.) Often animals can use either proximal or distal cues, depending on which are available. A now-classic demonstration was devised by Morris (1981; see Figure 8.4). A rat is placed in a circular pool of water in which it swims until it finds a small dry platform, a plexiglas cylinder standing somewhere in the pool. For some rats, the cylinder is black and visible above the water. Thus the platform can function as a beacon, and because rats would rather be dry than swim, they soon learn to approach it wherever it is in the pool. For other rats, the water is made opaque by the addition of milk, and the platform is transparent and slightly below the water surface. These rats must use distal cues, objects in the room surrounding the pool, to find the platform, and they also quickly learn to approach it, provided it stays in the same place from trial to trial. When the platform is removed on test trials, these rats still head directly to the correct location and swim around it as if searching for the platform (see Figure 8.4). This behavior has typically been taken as evidence for

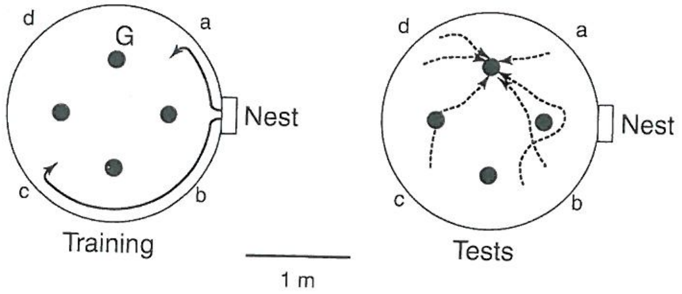

Figure 8.3. Hamsters use long-term memory of a location found by dead reckoning. In training a hamster was repeatedly lured from the nest around the edge of the arena along the two paths shown and then found its own way to the one baited cylinder (G) in darkness. In the tests animals were lured by each of the two possible paths from the nest to each of the four release sites, a–d. Subsequent paths of one hamster to the goal are shown. The paths from the familiar release sites, a and c, are only from trials with the novel path from the nest. After Etienne et al. (1998) with permission.

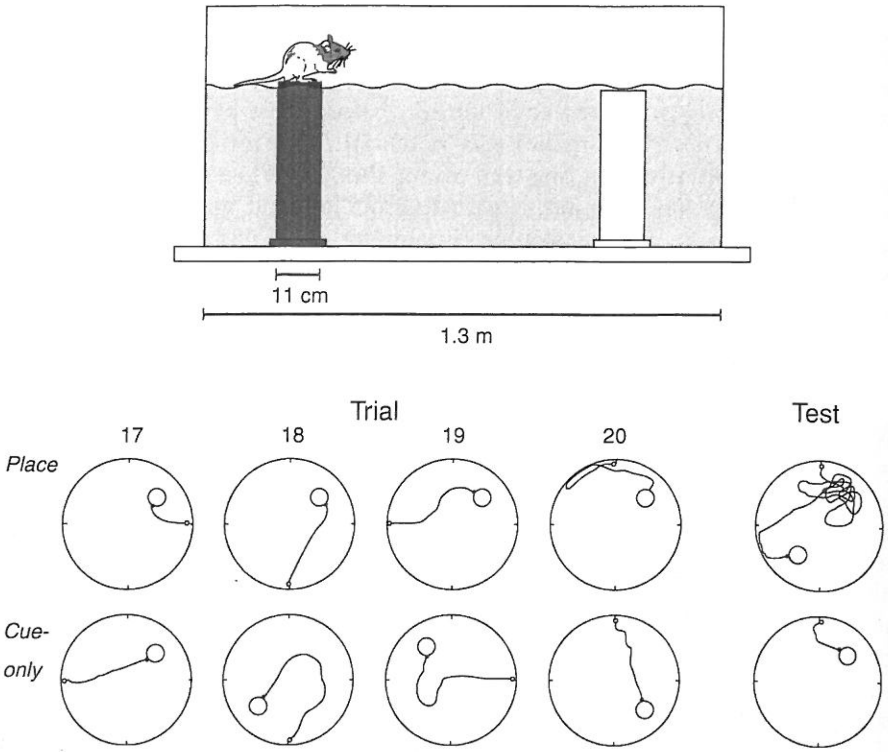

Figure 8.4. The Morris water escape task ("water maze"). At top, a cross section of the pool with a black visible platform and a white platform designed to be invisible to a swimming rat. Bottom: performance on trials 17–20 and a single test trial of one rat trained with the invisible platform always in the same place until the test ("place" condition) and one rat trained with the visible platform in a new place on each trial ("cue only" condition). After Morris (1981) with permission.

learning the specific place where the platform is, but it may often reflect instead learning what direction to head relative to distal cues (Hamilton et al. 2008).

Information from beacons is not inherently spatial because it is not vector information but rather information about value. Cues from a desired object or place, almost by definition, draw the animal to them. A classic subject in ethology is the analysis of simple mechanisms which bring this about (Fraenkel and Gunn 1961). Learned as well as unlearned features attract the animal to the goal: a fundamental effect of conditioning (Chapter 4) is that animals approach CSs associated with positive USs. For mammals, the intuition that beacons and landmarks demand different kinds of cognitive processing is supported by evidence from behavioral neuroscience (N. White and McDonald 2002). Rats with hippocampal lesions can still learn to approach a beacon like the dry platform in the swimming task, but they cannot learn tasks in which a goal is identified only by its spatial relationship to landmarks. But while finding a goal by approaching cues attached to it may be computationally simple, it has a major practical drawback: the animal must stay within range of those cues. In most natural environments, an animal that had to be able to see, smell, or hear its nest or possible food sources at all times would have its travels severely limited.

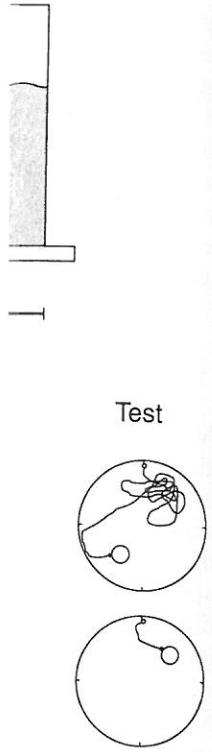

ɛction of the pool with a
imming rat. Bottom:
ıe invisible platform
ed with the visible
1981) with permission.

often reflect instead
ı et al. 2008).
t is not vector infor-
red object or place,
ct in ethology is the
el and Gunn 1961).
goal: a fundamental
CSs associated with
landmarks demand
ıce from behavioral
ocampal lesions can
ıming task, but they
atial relationship to
ached to it may be
ıe animal must stay
ıimal that had to be
imes would have its

8.1.3 Landmarks

When features of a goal are not immediately perceptible from a distance, other objects in fixed locations, that is, *landmarks,* can guide the animal to it. A classic demonstration of landmark use is Tinbergen's (1932/1972) study of homing in the digger wasp (*Philanthus triangulum*). These wasps lay their eggs in a number of burrows, which they provision with bees. Each bee that a wasp collects requires a separate foraging trip, so the female wasp has to learn the location of each of her burrows. This learning takes place during a brief orientation flight. When leaving the nest for the first time, the wasp turns and faces the nest entrance and flies around in ever-increasing loops, apparently inspecting the entrance and the objects around it (Figure 8.5a). If the objects surrounding an established nest are altered while the wasp is inside, a new orientation flight will be elicited the next time she departs (T. Collett and Lehrer 1993; Lehrer 1993).

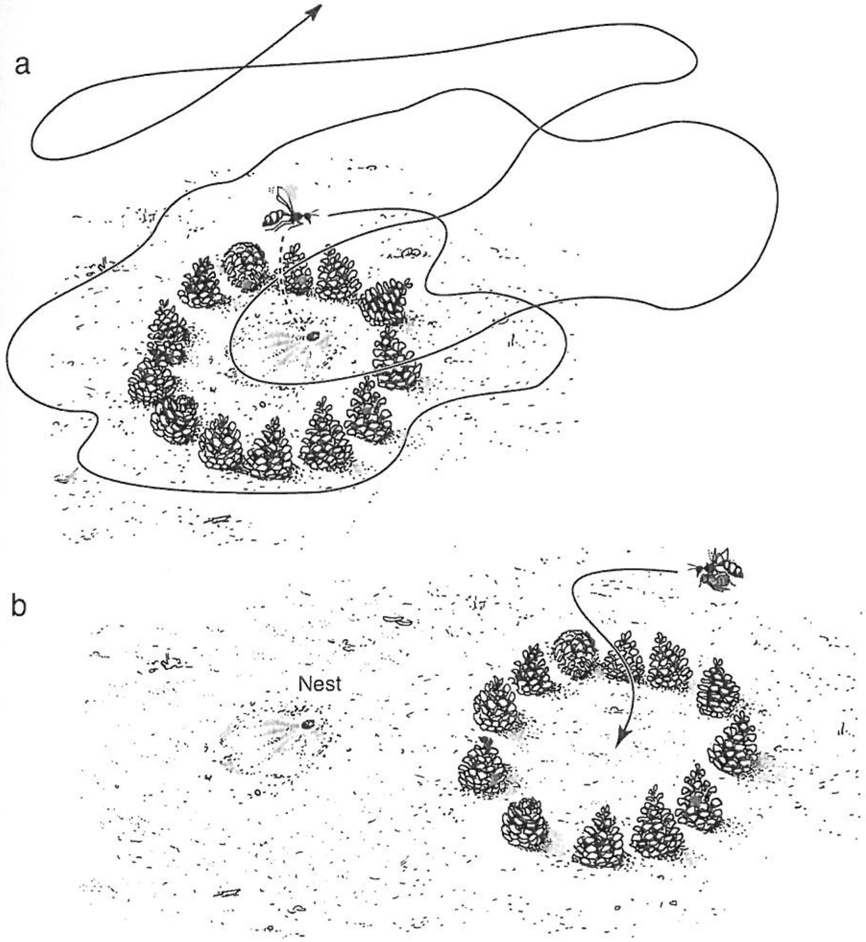

Figure 8.5. Control of orientation in the digger wasp (*Philanthus triangulum*) by nearby landmarks, a circle of pine cones. In a the wasp is shown making an orientation circle over the nest entrance before departing. After N. Tinbergen (1951) with permission.

To discover whether digger wasps were locating their nests using nearby landmarks, Tinbergen made a circle of pinecones around a nest while the wasp was inside and allowed it a number of trips in which to learn about them. Then they moved the pinecone circle to one side of the nest while the wasp was out foraging (Figure 8.5b). Although the nest entrance was still visible, returning wasps nearly always landed in the pinecone circle and searched for the nest entrance there. Only when the experimenters moved the pinecones back did she reenteer the nest. To discover which nearby landmarks the wasps learned about, Tinbergen and Kruyt (1938/1972) made landmark circles from two kinds of objects and tested the wasps with separate circles of each kind, one on each side of the nest. Wasps preferentially used as landmarks objects that were large, nearby, and three-dimensional. Such a preference makes functional sense. Large three-dimensional objects are more likely to be visible from a distance than small flat ones, and if perception of distances and directions obeys Weber's Law (Chapter 3), objects close to a goal localize it more accurately than objects farther away. Thus it is not surprising that similar preferences have been found in other animals, including European jays (*Garrulus glandularius*) (Bennett 1993) and honeybees (Cheng et al. 1987). Mechanistically, they likely reflect overshadowing during landmark learning (see Section 8.3). A landmark at a given distance supports more accurate localization the nearer it was to the goal in a training array composed of several landmarks (Goodyear and Kamil 2004).

How are landmarks used? Template matching and local views

One way to compute how to move toward a goal is to compare one's current view of the surroundings with a "snapshot" stored in memory of how the world looks from the goal. Honeybees appear to use such a mechanism. Bees were trained to find sugar water in a particular location in a laboratory room and tested with the familiar landmark array expanded or contracted. When a single landmark defining the goal's location was doubled in size, bees searched twice as far away from it as usual, that is, at the distance where the landmark would look the same as from the goal; conversely, when the landmark was half as big, bees halved the distance at which they searched (Cartwright and Collett 1983). The bee makes the matching task easier for itself by facing important landmarks in a standard compass direction, which it gets from its magnetic sense (T. Collett and Baron 1994). The animal apparently does not need to memorize how the goal looks from all directions. Chickens apparently behave similarly (Dawkins and Woodington 2000).

Figure 8.6 depicts a demonstration (Stürzl et al. 2008) that image-matching can be used to find a goal in a simple laboratory task. Food is buried in one corner of a rectangular enclosure with three black walls and one white one (panel a). Panoramic (i.e., 360°) images centered roughly at the intersection of wall and floor and taking in 115° vertically are recorded at the goal (figure 8.6b) and at other points throughout the arena. Computing the total pixel-by-pixel difference between the image at any location and the image at the goal gives a map of the arena indicating which way the creature relying on such images should move from each point to maximally reduce the difference between the current and the desired image (Figure 8.6c.) The arrows from most starting positions converge on the goal but there will be a substantial number of erroneous choices of the diagonally opposite corner, the "geometric errors" discussed in Section 8.1.5.

nests using nearby land-
est while the wasp was
about them. Then they
wasp was out foraging
, returning wasps nearly
est entrance there. Only
he reenteer the nest. To
it, Tinbergen and Kruyt
cts and tested the wasps
e nest. Wasps preferen-
and three-dimensional.
dimensional objects are
ones, and if perception
), objects close to a goal
it is not surprising that
ncluding European jays
(Cheng et al. 1987).
landmark learning (see
re accurate localization
d of several landmarks

compare one's current
mory of how the world
mechanism. Bees were
a laboratory room and
tracted. When a single
, bees searched twice as
e landmark would look
k was half as big, bees
and Collett 1983). The
portant landmarks in a
ic sense (T. Collett and
nemorize how the goal
imilarly (Dawkins and

t image-matching can be
ried in one corner of a
ne (panel a). Panoramic
l and floor and taking in
other points throughout
tween the image at any
ndicating which way the
to maximally reduce the
8.6c.) The arrows from
a substantial number of
metric errors" discussed

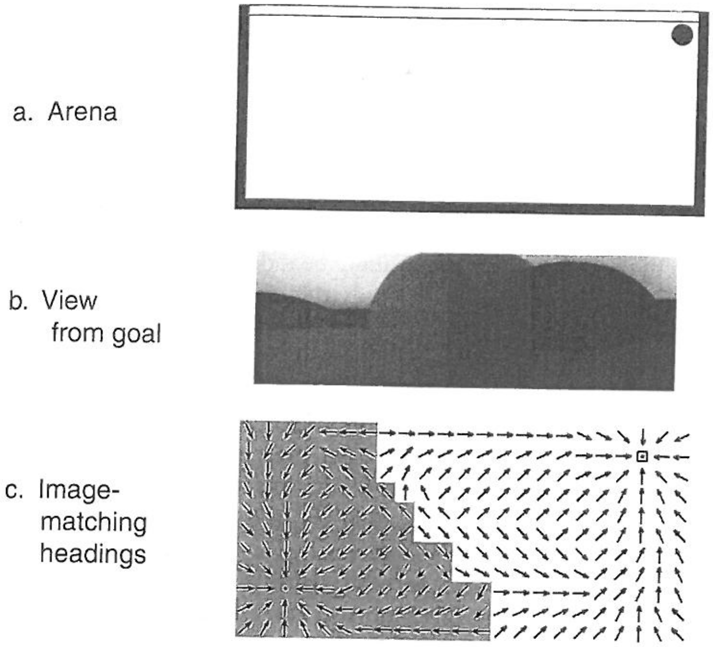
a. Arena

b. View from goal

c. Image-matching headings

Figure 8.6. a. Rectangular arena with one white and three black walls. Black dot indicates the location of buried food (the goal). b. Panoramic image of the arena as seen from the goal (the 360° view is unwrapped with the goal corner to right of center). c. Map of predicted headings for a creature moving at each point to maximally reduce the discrepancy between the current view and that at the goal. Note that from the majority of locations a creature following the arrows will arrive at the correct corner or its geometric equivalent. After Stürzl et al. (2008) with permission.

Figure 8.6 depicts a situation in which the animal is always within sight of the goal, but animals need to get close enough to the goal to use nearby landmarks in the first place. In principle this could also be accomplished by image matching. For example, a bee could have an "album of snapshots" (Cartwright and Collett 1987) from different locations within familiar terrain, each associated with a vector from that location to the hive. In rodents, this kind of mechanism is known as the *local view hypothesis* (Leonard and McNaughton 1990) "A location is nothing more than a set or constellation of sensory/perceptual experiences, joined to others by specific movements." (Leonard and McNaughton 1990, 366; see also McNaughton, Knierim, and Wilson 1995). Navigation based on learned links between local views is in effect what goes on in experiments in which people "move around" in a virtual environment by moving a joystick to reveal sequences of views simulating what one would see when moving around a neighborhood. With experience in realistic and complex virtual environments, people can plan novel routes using the same brain areas involved in "real" navigation (Hartley, King, and Burgess 2004).

Tinbergen's wasps must have used features of the terrain beyond the nest to find their way to within sight of the pinecone circles, but if an animal encounters similar landmarks or local views in different parts of its territory, it has to know which one is which. This problem can be solved by spatial context learning or occasion setting. For example, honeybees use distant landmarks or memory of the recent route to recognize ambiguous nearby landmarks. Bees were trained to find artificial nectar in each of two small featureless huts. Within each hut, the position of the food was specified by an identical array of four landmarks, but it was on the left of the landmarks in one hut and on the right in the other. The bees learned to search in the appropriate position, apparently remembering the global spatial context (Collett and Kelber 1988).

How are landmarks used? The vector sum model

Rather than using a whole visual panorama, animals may encode information about individual landmarks. But although a single beacon is sufficient to localize a goal, a single symmetrical landmark indicates only the distance to the goal. Without directional information, it can do no better than search in a ring around the landmark. Two discriminably different landmarks unambiguously specify a single position, and an array of three of more landmarks provides redundant information. To discover how information from such multiple landmarks is combined, animals can be trained to find a goal with two or more landmarks present, and then one or more of the landmarks is moved, in a so-called transformation test (Cheng and Spetch 1998).

Sometimes animals behave as if learning about only one of several available landmarks. For example, when gerbils were trained to search between two landmarks which were then moved further apart, the gerbils concentrated their searching in two spots, each at the correct distance and direction from one of the landmarks (T. Collett, Cartwright, and Smith 1986). In contrast, pigeons trained to search in a constant location in front of a wide stripe on the wall of a large rectangular box behaved as if averaging information from the conspicuous stripe and other features of the box (Cheng 1989). When the single landmark was shifted along the wall of the box in unrewarded test trials, the position where the birds pecked most shifted along with it, but typically not as much, that is, the birds averaged information from the landmark with some other feature, possibly the corners or visible features of the room outside the box (Figure 8.7). If the landmark was moved perpendicular to the wall of the box, searching shifted toward or away from the wall, but not as much as when the landmark was moved the same distance sideways. The nearby wall of the box seemed to be weighted relatively heavily in the bird's determination of how close to the wall to search. Black-capped chickadees (Cheng and Sherry 1992) and Clarks' nutcrackers (Gould-Beierle and Kamil 1996) also behaved similarly on comparable tests.

How are landmarks used? The multiple bearings model

But what exactly is being averaged? Are whole vectors averaged or are distances and directions computed separately? Cheng (1994) found some evidence that pigeons behave as if separately computing distance and direction from a single landmark. In a natural situation with landmarks more distant than features in a typical laboratory room, directional (or *bearing*) information by itself can be used to localize a goal surprisingly precisely, as illustrated in Figure 8.8. Bearing from a landmark to a goal, as in "the big pine tree is 40° northwest of my nest," does not change with distance, whereas judgment of goal-landmark distance, following Weber's Law, is less precise for more distant objects. Moreover, even if bearings are remembered with slight error, a goal surrounded by multiple landmarks, even quite distant ones, can be localized to the small area where the remembered bearings intersect (Figure 8.8). If animals' spatial judgments reflect these properties of the world, a number of predictions follow (Kamil and Cheng 2001). For instance, when an animal has learned to find a goal that is at a certain relative position, such as in the middle, between two

he food was specified by
the landmarks in one hut
he appropriate position,
ınd Kelber 1988).

code information about
ent to localize a goal, a
he goal. Without direc-
; around the landmark.
fy a single position, and
formation. To discover
animals can be trained
en one or more of the
g and Spetch 1998).
several available land-
tween two landmarks
ited their searching in
one of the landmarks
trained to search in a
large rectangular box
e and other features of
along the wall of the
ted most shifted along
information from the
e features of the room
dicular to the wall of
t as much as when the
all of the box seemed
how close to the wall
d Clarks' nutcrackers
parable tests.

or are distances and
idence that pigeons
ingle landmark. In a
a typical laboratory
d to localize a goal
landmark to a goal,
ange with distance,
Law, is less precise
mbered with slight
stant ones, can be
sect (Figure 8.8). If
number of predic-
mal has learned to
ddle, between two

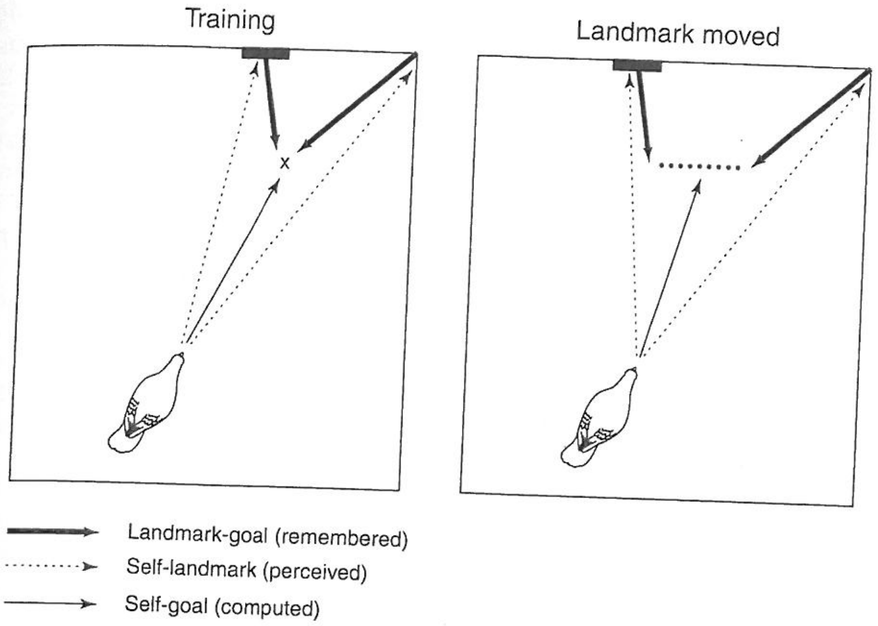

Figure 8.7. Hypothetical vectors involved in computation of the distance and direction to a goal (x) during training with a conspicuous landmark (black bar). The corner of the search space is treated as a second landmark. The self to landmark and landmark to goal vectors sum to produce the self to goal vector (the distance and direction resulting from summing two vectors is found by placing them head to tail). When the landmark is moved the animal will search somewhere along the dotted line, searching further toward the left the more heavily the black bar landmark is weighted relative to the corner of the box.

landmarks that vary in separation, direction errors should increase more slowly with interlandmark distance than distance errors.

Clark's nutcrackers are a particularly good species on which to test such predictions because they almost certainly need to rely on multiple and perhaps somewhat distant landmarks to relocate their buried caches under snow. Nutcrackers do behave as predicted by this multiple bearings hypothesis (Kamil and Cheng 2001) in several kinds of tests (Kamil and Jones 2000; Kamil and Goodyear 2001). Besides showing greater distance than direction errors, they more easily learn to locate a goal with a

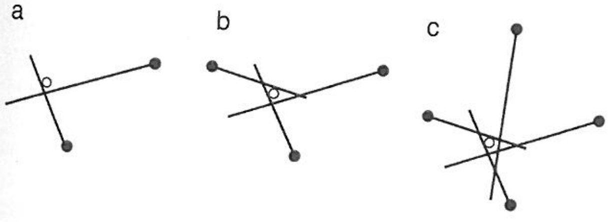

Figure 8.8. a. If an animal remembers only the compass directions (bearings) from the goal (open dot) to two landmarks and positions itself as near as possible to their intersection, even with small error in memory it can arrive reasonably close to the goal. b, c. Using more than two bearings confines search in a smaller area. After Kamil and Cheng (2001) with permission.

constant bearing to two landmarks than one at a constant distance from the line joining them. Pigeons are much less accurate than nutcrackers in laboratory tasks requiring them to use landmarks to search for buried seeds and do not show a clear difference between use of bearings and distances (Jones et al. 2002; Spetch et al. 2003). Although it is always difficult to be sure the training conditions are equated across species in such studies (but see Jones et al., 2002), the fact that this pattern of results has been found in more than one task and laboratory suggests that the nutcrackers have not only the exceptional spatial memory documented in Chapter 7 but exceptional ability at spatial localization.

A related finding that at first appeared to reflect a species difference in use of landmarks turned out instead to reflect differences in training methods. Spetch and colleagues (Spetch, Cheng, and MacDonald 1996; Spetch et al. 1997) found that pigeons trained to find the middle between two landmarks or in a square array of four landmarks in an arena or on a touchscreen behaved as if using only one landmark. When the landmarks were moved further apart in unrewarded tests, birds searched at the training distance from one of them. People behave in such tests as if they had learned "find the middle" (Figure 8.9). In this context, a report (Kamil and Jones 1997) that Clark's nutcrackers also behave as if learning a concept of middle might seem yet further evidence of the tendency of corvids to abstract concepts rather than memorize specific visual patterns as pigeons do (Mackintosh 1988). However, the pigeons in Spetch and colleagues' studies were trained with only a single interlandmark distance, whereas the nutcrackers were trained with multiple distances between the landmarks. The procedure used for the nutcrackers would be expected to teach the birds to weight both landmarks equally in determining distance, whereas relying on just one is a workable strategy for landmarks that never move. As this discussion predicts, when pigeons were trained like the nutcrackers with a variety of interlandmark distances, they also searched in the middle in the tests (Jones. et al. 2002).

8.1.4 Routes

"The animal got home because it had learned a route." As an explanation of accurate orientation, this statement is not very useful because "learning a route" can mean two different things. On the one hand, "learning a route" can refer to a mechanism of egocentric orientation in which an animal records the movements it makes in traveling between two places. This is usually referred to as *response learning* in psychology, to distinguish it from *place learning,* that is, use of landmarks. In the 1950s, considerable effort was devoted to testing whether rats learned mazes primarily as chains of responses or whether they learned about the relationships among places. Clark Hull is usually identified with the first view, and E. C. Tolman with the second. Like many controversies in psychology, this one was resolved—insofar as it ever was—by accepting that the answer to the question, "What does a rat learn in a maze?" is "It depends." Some conditions favor place learning and others, response learning (Restle 1957). Moreover, sometimes place and response learning go on in parallel and either one is used as the situation requires (Section 8.3.3).

A classic example of response learning comes from Konrad Lorenz's (1952, 109) depiction of how his pet water shrews followed their

> path-habits, as strictly bound to them as a railway engine to its tracks and as unable to deviate from them by even a few centimetres.... The shrews, running along the wall, were accustomed to jump on and off the stones which lay right in their path. If I

t distance from the line
kers in laboratory tasks
and do not show a clear
t al. 2002; Spetch et al.
conditions are equated
fact that this pattern of
atory suggests that the
documented in Chapter

cies difference in use of
ng methods. Spetch and
et al. 1997) found that
in a square array of four
ing only one landmark.
d tests, birds searched at
uch tests as if they had
eport (Kamil and Jones
oncept of middle might
act concepts rather than
h 1988). However, the
only a single interland-
ltiple distances between
ld be expected to teach
istance, whereas relying
nove. As this discussion
h a variety of interland-
Jones. et al. 2002).

explanation of accurate
a route" can mean two
fer to a mechanism of
nts it makes in traveling
rning in psychology, to
the 1950s, considerable
primarily as chains of
ong places. Clark Hull
the second. Like many
it ever was—by accept-
maze?" is "It depends."
earning (Restle 1957).
rallel and either one is

d Lorenz's (1952, 109)

tracks and as unable to
unning along the wall,
ight in their path. If I

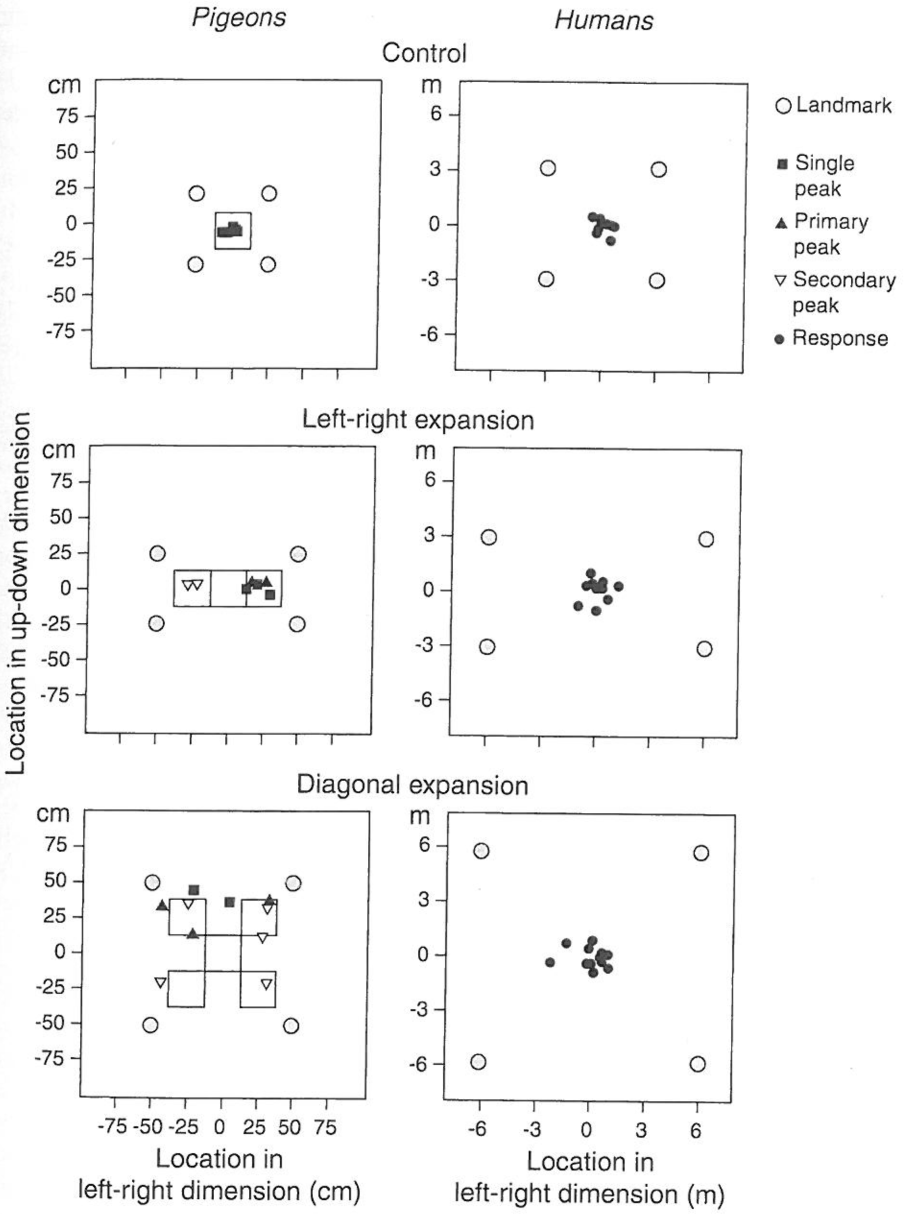

Figure 8.9. Setup and results of experiments testing how pigeons and humans use landmarks when trained to find the middle of a square array of landmarks in an open field. Redrawn from Spetch et al. (1997) with permission.

moved the stones out of the runway, . . . the shrews would jump right up into the air in the place where the stone should have been; they came down with a jarring bump, were obviously disconcerted and started whiskering cautiously left and right, just as they behaved in an unknown environment.

Gallistel (1990, 96–98) reviews analogous examples from the behavior of rats in mazes. As he points out, the animal must be keeping track of its distance and direction from the starting point (otherwise, it would not know where to jump), and it must use other cues to orient itself at the start. For the nearly blind water shrew, these must be tactile and/or olfactory cues gained by "whiskering." The disadvantage of sacrificing continuous monitoring of the environment for speed is that changes in the environment are not detected immediately. However, as Lorenz (1952, 111) pointed out, the shrew's brand of route learning has some advantages. It

> compensates the shrew for being nearly blind and enables it to run exceedingly fast without wasting a minute on orientation. On the other hand, it may, under unusual circumstances, lead the shrew to destruction . . . water shrews have broken their necks by jumping into a pond which had been recently drained. In spite of the possibility of such mishaps, it would be short-sighted if one were to stigmatize the water shrew as stupid because it solves the spatial problems of its daily life in quite a different way from man . . . by learning by heart every possible spatial contingency that may arise in a given territory.

In discussions of orientation in natural environments, route learning often refers to reaching a goal using a series of landmarks, that is, a series of stimulus-response (S-R) associations. This kind of orientation can be illustrated with examples of guides for hikers (O'Keefe and Nadel 1978). A person may be instructed "after crossing the bridge, turn left and proceed along the bank of the stream until you reach a hedge. Turn right and climb the hill." Similarly, an animal may learn its way around familiar territory by memorizing distances and directions of travel with respect to landmarks. When homing pigeons are repeatedly released from the same location a few kilometers from their loft, individuals adopt different routes, but each one takes the same route time after time (Biro, Meade, and Guilford 2004; but see Wiltschko, Schiffner, and Siegmund 2007). When honeybees (F. Dyer 1994, see Section 8.3) and desert ants (T. Collett and Collett 2004) repeatedly visit the same foraging site they too learn routes with respect to landmarks in addition to using path integration (see Section 8.2.3).

8.1.5 Environmental geometry

In 1986, Ken Cheng published a remarkable discovery. He had devised a simple test of spatial working memory in which rats found food in a large rectangular box placed within a dark room, were removed from the box for about a minute and then replaced in an identical box differently oriented in the room to dig for the now-buried food. In test trials, no food was present and digging was recorded. The rats showed good memory for locations of food which they had experienced just once, in that they dug in the correct place at above chance levels. But amazingly, they dug nearly as often at the diagonally opposite point in the box like the hypothetical view-matching creature in Figure 8.6. Notice that in diagonally opposite locations the animal's relationship to the box's geometry is the same. For example, a short wall may be on the left, a long wall on the right. And some correlate of geometry, the box's shape, seems to be what the rats are paying most attention to (Cheng 1986; Chapter 6 in Gallistel 1990). For if geometrically identical locations are made more discriminable, for instance by coloring one long wall white and the others black as in the enclosure depicted in Figure 8.6, the rats still make diagonal errors. Similarly, placing distinctive panels with different patterns and odors in the corners still does not eliminate the tendency to make primarily diagonal errors.

t know where to jump),
r the nearly blind water
d by "whiskering." The
environment for speed
nediately. However, as
oute learning has some

t to run exceedingly fast
d, it may, under unusual
have broken their necks
spite of the possibility of
atize the water shrew as
in quite a different way
ngency that may arise in

ite learning often refers
es of stimulus-response
ated with examples of
ay be instructed "after
of the stream until you
a animal may learn its
nd directions of travel
edly released from the
adopt different routes,
Meade, and Guilford
7). When honeybees
ett and Collett 2004)
with respect to land-
.3).

devised a simple test of
ectangular box placed
nute and then replaced
e now-buried food. In
he rats showed good
once, in that they dug
dug nearly as often at
v-matching creature in
al's relationship to the
the left, a long wall on
ns to be what the rats
allistel 1990). For if
r instance by coloring
cted in Figure 8.6, the
panels with different
ncy to make primarily

Cheng took pains to force his rats to rely on spatial cues within the boxes. They were in a dark and relatively featureless room. Testing the rat in a different box differently oriented in the room meant it could not rely on dead reckoning to return it to the same location in space after it had been removed from the first box. These conditions are crucial for control by geometry. When Cheng's experiments are repeated but with the room visible outside the box and the test and exposure boxes in the same location within the room, rats search almost exclusively in the correct location and make no more diagonal than other kinds of errors. When they are disoriented by making the room dark, not always having the exposure and test boxes in the same place, and being gently rotated between exposure and test, the same rats make as many diagonal errors as correct responses (Margules and Gallistel 1988, Experiment 3).

Cheng's (1986) findings turn out to have remarkable generality across vertebrates. Young chickens, pigeons, black-capped chickadees, two species of fish, and monkeys also encode the locations of goals relative to the geometry of an enclosure, even in the presence of features like corner panels or a colored wall that disambiguate the geometry (Cheng and Newcombe 2005). Like rats (Wall et al. 2004), these animals can eventually perform well in a reference memory task with food in the same place on every trial relative to such features. But even when a feature is the best cue to the goal, they still learn the relationship of the goal to the box's geometry, as shown by searching in geometrically correct locations when the features are removed. Geometry even takes precedence over featural information when young children are tested similarly to rats in a working memory task. Hermer and Spelke (1994) showed college students and 20-month-old toddlers the location of an object in a room and then asked them to find it after they had shut their eyes and turned themselves around ten times. If the room was white and featureless, the students and the toddlers behaved just like Cheng's rats—not surprisingly, since they had no cues to disambiguate the correct corner from its diagonal. When the room was given one blue wall, the students searched mostly in the correct place, but the toddlers were just as confused as before. Like Cheng's rats, they could be provided with salient features in the room (a teddy bear, a toy truck) that they could use for orientation, but when they were disoriented by being rotated before searching, they still fell back on purely geometric information.

By the time children are about six, they use featural cues as adults do (Cheng and Newcombe 2005). Moreover, when adults' attention is occupied with a second cognitive task during the retention interval in a test like Hermer and Spelke's, they fall back once more on geometry (Hermer-Vazquez, Spelke, and Katsnelson 1999). In Cheng's (1986) original discussion, the fact that shape of the environment seems to take priority over features of the very surfaces that define that shape was interpreted as meaning that environmental geometry is processed in a dedicated cognitive module, impenetrable to other spatial information. On this view, developmental changes in use of geometry show that although humans share the geometric module with other vertebrates, language allows them to overcome its limitations (Hermer and Spelke 1994; Wang and Spelke 2002). But the claim that language is critical here is controversial (Cheng and Newcombe 2005; Newcombe and Ratliff 2007). For example, the importance of featural cues relative to geometry depends on the size of the enclosure. Children (Learmonth et al. 2008) as well as chicks and fish are more likely to use features in a relatively large space, in some cases possibly because the features are simply larger (Chiandetti and Vallortigara 2008; N. Y. Miller 2009). Thus even if—as discussed further in the next section—most vertebrates have a geometric module, its output may be combined in adaptive ways with other information (Newcombe and Ratliff 2007).

Of course not only enclosures but also configurations of landmarks have a shape, like the linear and square arrays of landmarks in Figure 8.9. However, the animals were not disoriented in those studies so they are not strictly comparable to the studies of enclosure geometry being discussed here. The limited evidence available indicates that disoriented rats and people do not encode the shape of an array of objects as such (Wang and Spelke 2000; Skov-Rackette and Shettleworth 2005). So why should the global shape of the surrounding environment be so important, and what about it are animals encoding anyway? One answer to the first question is that sensitivity to overall geometry is a mechanism for reorientation, or getting a heading (Wang and Spelke 2002). A not uncommon experience of disorientation and reorientation occurs when one emerges from an unfamiliar subway exit into the street and does not at first know which way is which. The claim is that the overall shape of the surroundings permits reorientation, after which specific environmental features can be identified.

What is geometry?

What it is about shape that is encoded is still unclear. In a rectangular enclosure a rat could encode its position relative to a box's geometry as a certain distance from a corner with a long wall on the right and a short wall on the left. That is, it might encode comparatively local spatial information about absolute or relative (Kelly and Spetch 2001) wall lengths and their left-right position or *sense* and perhaps also the angle at which they meet (Tommasi and Polli 2004). In contrast, using more global spatial information, the animal might extract the principal axes of the space and locate the goal relative to them, for example, at one end of the long axis and to the right (Figure 8.10). (In a symmetrical shape like a rectangle, the long axis is simply the line that divides it in half lengthwise.) Testing what is used requires transforming the space in some relevant way once the animal has learned to use geometric cues.

In one such test, Pearce, Good, Jones, and McGregor (2004) trained rats to find the dry platform in one corner of a rectangular water tank and then gave them unrewarded tests in a kite-shaped tank made by taking the rectangle apart along one diagonal, flipping one of the resulting halves over and putting the enclosure back together (see figure 8.10). Thus it now had two right-angled corners, only one of which had the same

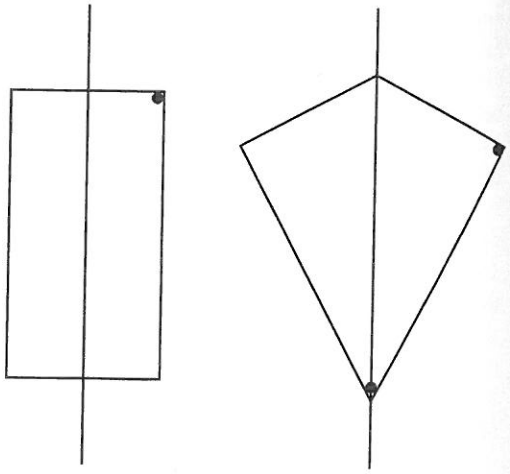

Figure 8.10. Layouts of the training (left) and testing enclosures in the experiment by Pearce et al. (2004) described in the text, showing principal axes (long vertical lines). Black dots in the kite-shaped arena indicate where rats spent most time searching for the platform in trials without the platform after training to go to the corner of the rectangle with the back dot. After Cheng and Gallistel (2005) with permission.

vith other information

marks have a shape, like
er, the animals were not
the studies of enclosure
ndicates that disoriented
such (Wang and Spelke
the global shape of the
t are animals encoding
o overall geometry is a
d Spelke 2002). A not
curs when one emerges
irst know which way is
permits reorientation,

lar enclosure a rat could
ice from a corner with a
t encode comparatively
etch 2001) wall lengths
le at which they meet
nformation, the animal
l relative to them, for
In a symmetrical shape
lf lengthwise.) Testing
ay once the animal has

rained rats to find the
gave them unrewarded
t along one diagonal,
re back together (see
of which had the same

adjoining long and short walls with the same sense as the training corner. Much of the rats' time was spent searching for the platform in this corner, as if they had learned purely local cues. However, they searched about as much at the newly created sharp-angled corner, a finding better explained by the more global, principal axis, account (Cheng and Gallistel 2005). This latter account can also explain the results of an experiment in which young chicks were tested in transformations of a rhomboid-shaped enclosure (Tommasi and Polli 2004), although again the original authors favored a more local account of what their subjects had learned. However, local geometric features and global axes are not the only possibilities. The image matching mechanism depicted in Figure 8.6 does pretty well with several studies involving transformations of kite-shaped arenas even though it does not assume animals have any geometric information as such (Cheung et al. 2008). In summary, then, although the basic phenomena of geometry learning have proven remarkably robust, how best to account for them remains controversial (see Cheng 2008).

8.2 Modularity and integration

8.2.1 Spatial modules?

Section 8.1 makes clear that animals often have a wealth of cues for orientation available simultaneously—far and near landmarks, cues emanating from important goals (beacons), environmental shape, the *idiothetic* (self-generated or internal) cues used in path integration, memory for the chain of responses that got them from one place to another. These serve as input to distinct servomechanisms demanding different implicit computations. For instance, dead reckoning is a working memory process that takes as input some correlate of distance and direction traveled and outputs an approximation of the vector back to the starting place. The process revealed in geometry-learning experiments uses unknown parameters of the surrounding space to locate the animal relative to a global heading. Orienting by landmarks takes as input perceived self-landmark vectors and returns a vector from the current position to the remembered location of some goal. From a functional point of view then, spatial information processing consists of modular subprocesses.

But as discussed in Box 2.2, claims of modularity in the cognitive sciences are nearly always controversial because candidate modules seldom fit all Fodor's (1983) classic criteria. A debatable feature of possible spatial modules is the extent to which they are encapsulated, or impenetrable to anything other than their own specific kind of input (see e.g., Cheng and Newcombe 2005). What originally led Cheng (1986) and others (e.g., Gallistel 1990; Wang and Spelke 2002) to emphasize the modularity of spatial processing was not differences in implicit computations so much as striking observations of apparently stupid behavior in which one kind of spatial information is used to the exclusion of others that animals are manifestly sensitive to. The displaced desert ant runs right past its nest, the mother gerbil searches a blank wall even within range of the smells and cries of her babies, the water shrew jumps over a nonexistent stone, the rat turns its back on a conspicuous landmark that defines the correct corner and digs on the opposite side of the box. Such behavior suggests the animals are using one encapsulated module at a time. Indeed, in natural environments redundant cues are normally not dissociated, so relying on just one at a time is likely to work and—as with Lorenz's water shrews—may be more efficient than processing lots of cues at once. Reliance on one cue at a time may also reflect the path of evolution. More sophisticated and flexible orientation may have evolved by the addition of new modules rather than the modification of old ones. But in any case, under many circumstances animals equipped with multiple spatial modules or servomechanisms combine their outputs. In Fodorian terms, modular spatial mechanisms provide input to central decision making (Cheng 2005; Cheng and Newcombe 2005). In this section we consider ways in which multiple spatial inputs are combined to reach a decision about which way to go. A basic research strategy here is to place cues in conflict with one indicating one goal location and one, another. Does the animal search at one place, at the other, or somewhere in between? The relative weightings of different sources of information may change with the conditions. If the conflict between them is too great, animals appear to fall back on one and disregard the other. In some situations one set of cues is primary, providing a context in which other cues are used.

8.2.2 Bayesian averaging

In the vector sum model discussed in Section 8.1.3 information from two or more landmarks, that is, within one module, is averaged. However, although the example in Figure 8.7 indicates that some landmarks are weighted more heavily than others, the model does not specify how these weightings are determined. Functionally, more informative landmarks should be weighted more heavily. Elegant quantitative support for this supposition comes from human psychophysical studies investigating how two or more cues are weighted in determining perceptual localization. For example, in the *ventriloquist effect* people perceive the ventriloquist's voice coming from his puppet's mouth, as if the visual cue of a moving mouth overrides the binaural auditory cue to the location of the sound source. This phenomenon has been brought into the laboratory with stimuli consisting of a blob shown on a video screen simultaneously with a sound presented through stereo headphones, an event experienced as a ball hitting the screen (Alais and Burr 2004). Two such stimuli are briefly presented in succession, and people judge which is to the left. As one might expect, the more blurry the blob the greater the variance in judging its location when it is presented alone. More importantly, the more blurry the blob, the more combined blob+sound stimuli are localized toward the (virtual) sound source. In effect, subjects localize the bimodal stimulus at a weighted average of the locations of its components, weighting each component in inverse proportion to its variance. Such weighting on the basis of prior knowledge of probability distributions (here, "knowledge" is direct perception of fuzziness or sharpness) is prescribed by Bayes' law, according to which it is the optimal way to estimate any metric value. It applies widely in comparable situations (Cheng et al. 2007). Other aspects of Bayesian decision making are of broad interest in psychology (Chater, Tenenbaum and Yuille 2006), but they are beyond the scope of this book.

Although Bayes' Law provides quantitative functional predictions for weighting two or more information sources relevant to localizing a single goal, few data on animal landmark use are adequate to test it precisely because this requires data on the variance in judgments when each information source is presented alone (see Cheng et al. 2007). But a number of studies have provided data consistent with it. For example, on the reasonable assumption that distance judgments obey Weber's Law (i.e., their variance increases with the distance being judged), landmarks should be weighted less the further they are from a goal. An elegant illustration of this principle comes from a study of Clark's nutcrackers relocating their caches (Vander Wall 1982). Birds buried pine seeds throughout a 1.5 meter long oval arena with several prominent landmarks

es rather than the modifica-
ices animals equipped with
outputs. In Fodorian terms,
sion making (Cheng 2005;
ys in which multiple spatial
y to go. A basic research
one goal location and one,
or somewhere in between?
nay change with the condi-
ear to fall back on one and
ary, providing a context in

nation from two or more
ver, although the example
more heavily than others,
nined. Functionally, more
Elegant quantitative sup-
ical studies investigating
ceptual localization. For
triloquist's voice coming
uth overrides the binaural
omenon has been brought
own on a video screen
lphones, an event experi-
o such stimuli are briefly
. As one might expect, the
; its location when it is
lob, the more combined
source. In effect, subjects
locations of its compo-
s variance. Such weight-
ns (here, "knowledge" is
Bayes' law, according to
ie. It applies widely in
Bayesian decision mak-
n and Yuille 2006), but

redictions for weighting
ingle goal, few data on
this requires data on the
nted alone (see Cheng et
nt with it. For example,
Weber's Law (i.e., their
should be weighted less
principle comes from a
Vall 1982). Birds buried
l prominent landmarks

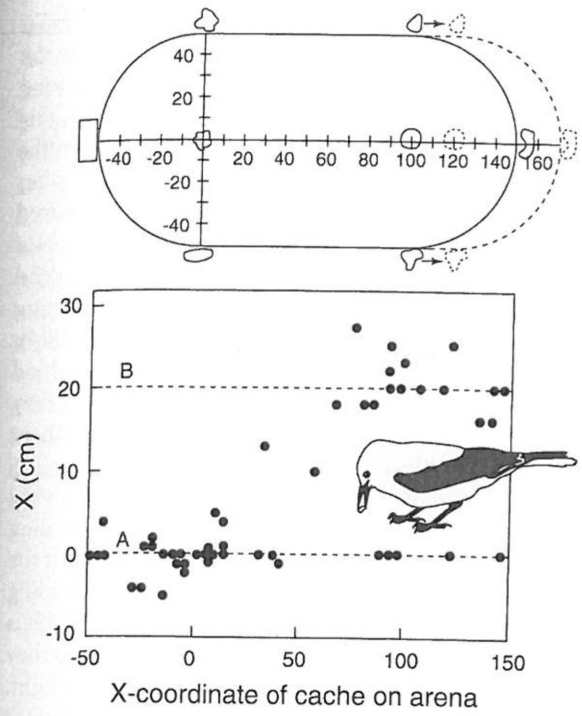

Figure 8.11. Setup and results of experiment to investigate response of Clark's nutcrackers to moved landmarks. Scale on diagram of the arena is in centimeters. Data are the distance between the location of the nutcrackers' probes for hidden seeds and the actual left-right position of the caches, as indicated on the map of the arena. Lines A and B represent, respectively, the loci of probes if the birds ignored the moved landmarks or followed them entirely. Redrawn from Vander Wall (1982) with permission.

at each end. The arena was then expanded by shifting all the landmarks at the right hand end 20 centimeters to the right (Figure 8.11). Thus caches near the right end of the arena were nearer to shifted landmarks than were caches on the left end. Birds probed farther from the stationary position of their caches and closer to a position shifted 20 centimeters the closer those caches were to the shifted landmarks. The graded effect shown in Figure 8.11, with searches in the middle of the arena shifted an intermediate distance, indicates that the moved landmarks were averaged with stationary ones, with landmarks close to a cache weighted more heavily than those further away.

Bayesian averaging should also apply when information from two spatial modules is being combined. One likely example comes from a study of honeybees in which directional information provided by a line of landmarks was put into conflict with direction given by the sun compass (Chittka and Geiger 1995). Many bees followed the landmark at displacements up to about 15°, but, as in some examples coming up next, they ignored the landmarks when they were moved too far. In Bayesian terms, averaging does not make sense if the possible positions indicated by the separate cues do not overlap because the prior probability that the goal is located between them is zero (further discussion in Cheng et al. 2007).

8.2.3 Parallel processing and hierarchical use

Rather than averaging the outputs of different spatial servomechanisms, animals may use them one at a time in a hierarchical manner. This often seems to be true when dead reckoning is involved. In numerous species and situations dead reckoning appears to be obligatory, always going on in the background and available as a backup when other cues fail, even when those cues were originally learned with reference to dead reckoning. One example comes from a study of rats by Whishaw and Tomie (1997). In a lighted room with plenty of landmarks, a rat's home cage was placed below the edge of a circular arena onto which the rat could climb and search for large food pellets which it carried back to the cage to eat. Because the cage was out of sight, rats initially had to use dead reckoning to return home, but with repeated trials from the same starting point they could learn to use landmarks as well. That this is what happened was confirmed by tests with a new starting point. In such tests, rats picked up the food and ran to the usual location of the cage with respect to landmarks. Not finding it there, they returned successfully to the new start location, using dead reckoning. Similarly, hamsters hoarding food immediately revert to dead reckoning when familiar landmarks are not visible (Etienne 2003; Etienne and Jeffery 2004). And as Figure 8.12 shows, when ants have made repeated trips home with food through a channel of a given length and are tested in shorter channels, after they emerge from the channel onto open ground they run in the direction that takes them home by dead reckoning (M. Collett et al. 1998).

As well as a backup, dead reckoning is used implicitly as a reference, in identifying landmarks in the first place (Cheng. et al. 2007). For example, hamsters hoarding food in the dark use a single small light as a landmark to return home. However, if the light is moved too far relative to the nest before the hamsters depart for the hoarding site, some of them ignore it and fall back on dead reckoning (Etienne. et al. 1990). It is as if dead reckoning leads the hamsters to expect the light in a certain position, so they disregard it when it is too far from that position, in effect treating it as a different light. Rats behave similarly (Shettleworth and Sutton 2005). Evidently a familiar landmark is recognized as such from its location with respect to the animal's internal position sense. Interestingly, if hamsters are repeatedly led astray by landmarks, they learn to rely more on dead reckoning and less on landmarks, as if recalibrating their relative weightings (Etienne 1992).

Other information about global spatial position may also determine how landmarks are used. For example, in a working memory task, black-capped chickadees learned the location of a single baited feeder in an array of four differently decorated feeders on the wall of an aviary and then searched for it later in a test of memory. When the array was moved along the wall for the test, they searched first in the feeder closest to the baited feeder's original position in the room (Brodbeck 1994). The birds

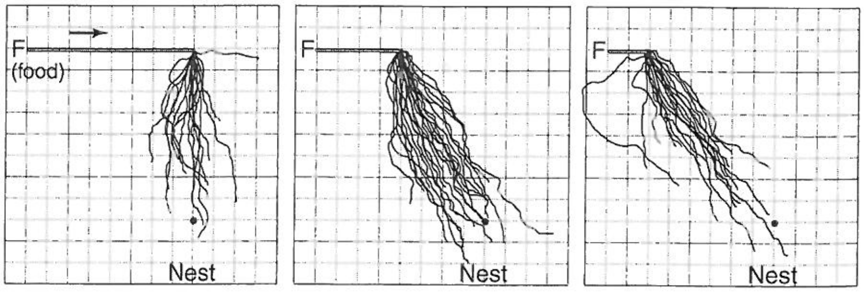

Figure 8.12. Trajectories of individual ants trained to find food at the end of an 8-meter-long channel and released at F in test channels of different lengths some distance away. Thus "nest" is the location where the nest would be relative to the beginning of the homeward trip at F in each diagram. From left to right, the test channel is 8, 4, or 2 meters long. Redrawn from Collett et al. (1998) with permission.

and Tomie (1997).
ıs placed below the
ırch for large food
as out of sight, rats
ated trials from the
That this is what
:h tests, rats picked
to landmarks. Not
cation, using dead
to dead reckoning
and Jeffery 2004).
's home with food
hannels, after they
on that takes them

ence, in identifying
hamsters hoarding
ne. However, if the
rt for the hoarding
ıe. et al. 1990). It is
in position, so they
as a different light.
familiar landmark
's internal position
arks, they learn to
ating their relative

ermine how land-
capped chickadees
ferently decorated
a test of memory.
d first in the feeder
k 1994). The birds

n 8-meter-long
ıy. Thus "nest" is the
at F in each diagram.
ett et al. (1998) with

---

tended to search next in the correct position in the array of feeders. However, when the array was moved too far along the wall, performance fell to chance, as if the birds did not recognize the feeders out of their global spatial context. In the same experiment, local cues such as color on the baited feeder were occasionally placed in conflict with spatial cues by swapping the formerly baited feeder with another feeder in the array for unrewarded probe trials (Figure 8.13). The chickadees went first to the feeder in the formerly baited location, even though it now looked different. Finding no peanut there, they tended to search next in the feeder with the correct color and pattern. Much of the time, these birds used the normally redundant cues hierarchically: global spatial, local array, and color/pattern. In contrast, dark-eyed juncos, which do not store food in the wild, weight color and pattern cues about equally with spatial cues (Brodbeck 1994). The same pattern of species difference is found in an analogous operant task (Brodbeck and Shettleworth 1995) and in two other pairs of food-storing versus non-storing birds (Clayton and Krebs 1994). Analogous tests dissociating local features from spatial cues have revealed differences between men

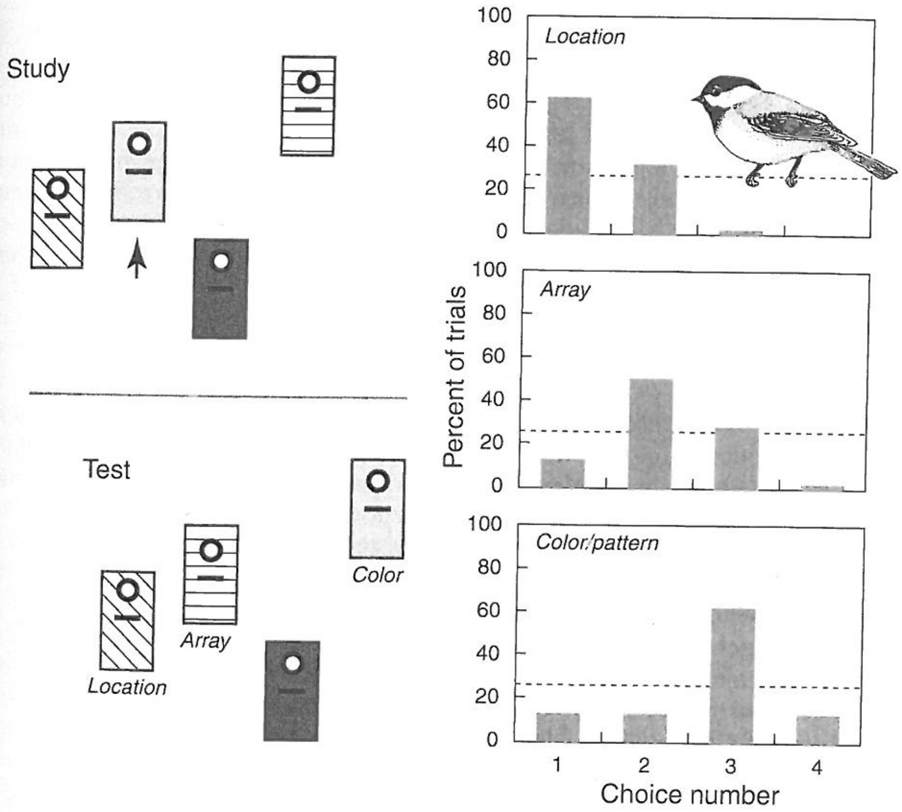

Figure 8.13. Setup and results of experiment investigating which cues to the location of reward are used by black-capped chickadees. On test trials the arrangement of feeders and their location in the room was changed as indicated between the study phase, in which the birds found and ate part of a peanut in one feeder, and the unrewarded test phase. Data are the proportion of trials in which the birds looked first, second, and so on, into the feeder in the formerly rewarded location in the room, in the same position in the array of feeders, or into the feeder of the rewarded color. Redrawn from Brodbeck (1994) with permission.

and women (Jones and Healy 2006) and between children and apes (Haun et al. 2006). Among birds, food storers' greater reliance on spatial cues may be related to their need for spatial memory in the wild (Chapter 7).

Using familiar cues in a hierarchical manner could result from learning some more strongly than others in the first place, because they overshadow other cues. For instance, perhaps chickadees use spatial cues before color cues when given a choice because they remember spatial cues better. However, although this kind of explanation may be correct for this case (Shettleworth and Westwood 2002), it is unlikely to apply in general. The examples involving dead reckoning sketched earlier in this section illustrate the more general principle that egocentric and allocentric orientation mechanisms operate in parallel (Burgess 2006). Possible parallel operation of multiple mechanisms is examined further in section 8.3.3.

8.3 Acquiring spatial knowledge: The conditions for learning

Most recent discussions of the conditions for spatial learning have been strongly influenced by O'Keefe and Nadel's (1978) claim that there is special spatial (*locale*) learning system, distinct from associative learning (the *taxon system*). The locale system is responsible for acquiring a cognitive map of the environment through exploration whereas the taxon system includes response learning, route learning, and classical conditioning, in effect all forms of associative learning. Exploration clearly does have an important role in spatial learning, as we see from research reviewed next. However, recent studies based on ideas about associative learning that largely postdate O'Keefe and Nadel's (1978) book support alternatives to the idea that all kinds of spatial information are spontaneously integrated into a unitary maplike representation.

8.3.1 Exploration

Exploration was a problem for S-R learning theory because it apparently resulted in learning without reinforcement, but in the 1960s the idea that behavior could be spontaneous and continue without reinforcement became more acceptable (Berlyne 1960; Hinde 1970a). The tendency to explore novel objects and environments is one of the best examples of special behaviors that expose animals to the conditions for learning. The rat sniffing a novel object, the young pigeon flying in circles over its loft, or the bee performing an orientation flight (Wei, Rafalko, and Dyer 2002) are actively exposing themselves to objects and spatial relationships that they need to learn about.

Spatial learning begins in the area around an animal's natal nest or burrow. A typical altricial rodent like a ground squirrel ventures out of its burrow a few weeks after birth but stays close to the entrance, maybe just rearing up and looking around from the mouth of the burrow. As the days pass, it makes longer and longer excursions around its mother's territory. Knowledge of the whole territory may be built up by connecting a series of "local charts," detailed knowledge about areas around important sites for food or refuge (see Figure 8.14; Poucet 1993). Indeed, one of the functions of territoriality may be to permit animals to acquire information that allows them to get around more safely and efficiently than they could in unfamiliar areas (Stamps 1995). For terrestrial animals, information from dead reckoning may be primary here, telling the animal where it is relative to its nest or burrow. By integrating the perceived

and apes (Haun et al.
cues may be related to

ım learning some more
adow other cues. For
s when given a choice
ı this kind of explana-
2002), it is unlikely to
:etched earlier in this
ıd allocentric orienta-
parallel operation of

ıing

g have been strongly
pecial spatial (*locale*)
ı system). The locale
:nvironment through
ning, route learning,
earning. Exploration
ve see from research
associative learning
rt alternatives to the
:grated into a unitary

pparently resulted in
t behavior could be
acceptable (Berlyne
environments is one
o the conditions for
n circles over its loft,
er 2002) are actively
need to learn about.
l nest or burrow. A
f its burrow a few
ıst rearing up and
ıss, it makes longer
edge of the whole
s," detailed knowl-
: (see Figure 8.14;
may be to permit
ıd more safely and
)5). For terrestrial
r here, telling the
ting the perceived

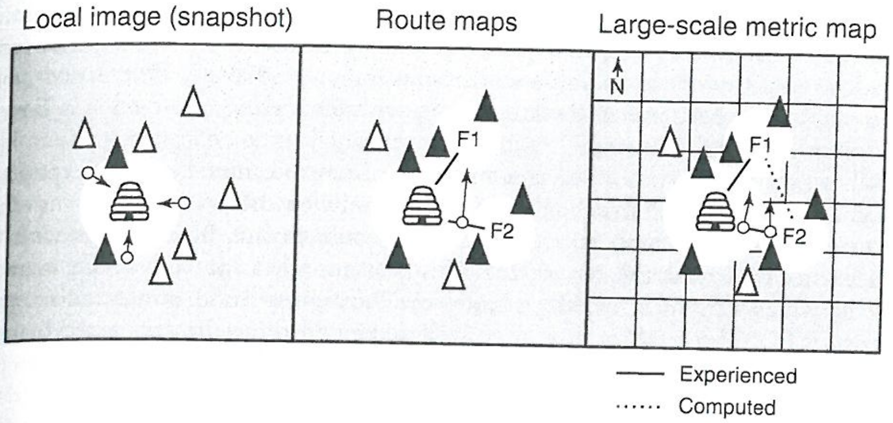

Figure 8.14. Three ways in which an animal may encode spatial information about the area around its home (the beehive). F1 and F2 are two feeding sites, the black triangles are known landmarks, and the white triangles are unknown ones. White areas are familiar to the animal; thin solid arrows are paths the animal (white dot) can take on the basis of the specified kind of representation. The grid in the large-scale map represents the idea that only in this kind of representation is information about different locations related to a common coordinate system, allowing the animal to compute a novel route (the dotted line). After F. Dyer (1996) with permission.

egocentric coordinates of prominent landmarks with this information, an animal can learn the position of nearby landmarks relative to its home (Gallistel and Cramer 1996; McNaughton et al. 1996).

The acquisition of spatial knowledge in the wild has been studied most in bees and homing pigeons. Pigeon racers have accumulated a vast fund of lore about what is necessary for the birds to learn the location of the loft (Keeton 1974; Wallraff 2005). Training racing pigeons typically begins by letting young birds fly around close to the home loft and then releasing them increasing distances away. In contrast, laboratory studies of exploration and spatial learning typically begin by dumping an animal into a completely novel environment. Even here the tendency gradually to venture further and further from a central place, presumably building up spatial knowledge, can be observed. For instance, rats placed in a large room to live travel over more and more of it in successive nights and gradually organize the space into nesting sites, food stores, runways, and latrines (Leonard and McNaughton 1990).

The two paradigms that have been used most extensively to study learning through exploration are habituation and tests of latent learning in mazes. In Chapters 5 and 7 we have seen how moving objects around, removing them, or introducing new ones elicits investigation of the altered object or location, evidence that the animals knew the features of the environment before it changed. This approach can also reveal what free-ranging animals know about their environment. For instance, wild rats ate less than usual from a familiar feeder displaced as little as a foot, showing they had learned its location quite precisely (Shorten 1954; see also Shillito 1963). And to take a rare example of spatial memory not involving food, free-ranging male thirteen-lined ground squirrels return to locations in their large (average 4.7 hectare) home ranges where they have previously encountered females. If the female has been removed, they spend longer searching for her if she had been about to go into heat the day before. The males also visit a female's territory earlier in the day when she is potentially receptive than otherwise, as if planning their route based on memory of the female's state (Schwagmeyer 1995). Male meadow voles show analogous behavior in the laboratory, with females encountered in a T-maze (Ferkin et al. 2008).

What aspects of exploration are important? Does the animal have to experience different routes through the environment, different views of it, or what? Some of the best examples of attempts to answer such questions come from studies with the Maier three-table task (Maier 1932a). This is essentially a spatial working memory task in which rats must rapidly encode the location of food in a familiar space (Figure 8.15). Three tables in a large well-decorated room are connected by a Y-shaped runway with a central platform. At the beginning of a trial, a rat is allowed to explore the whole apparatus, which is empty of food. It is then placed on the day's goal table with a large pile of food. After the rat has eaten for a few minutes but before it has depleted the food, it is placed on one of the other two tables, from which it may return and finish its meal. Trials are typically run only once a day, with the goal table changed from trial to trial. Experienced rats typically do quite well, even with delays of hours between feeding on a table and testing, but accurate choice of the goal table depends on prior opportunity to explore the maze (Maier 1932a; Stahl and Ellen 1974).

To discover whether rats can link together two parts of space they have never experienced closely together in time, Ellen, Sotere, and Wages (1984) restricted experience in the exploration phase. Three groups of rats had 15 minutes a day to explore the

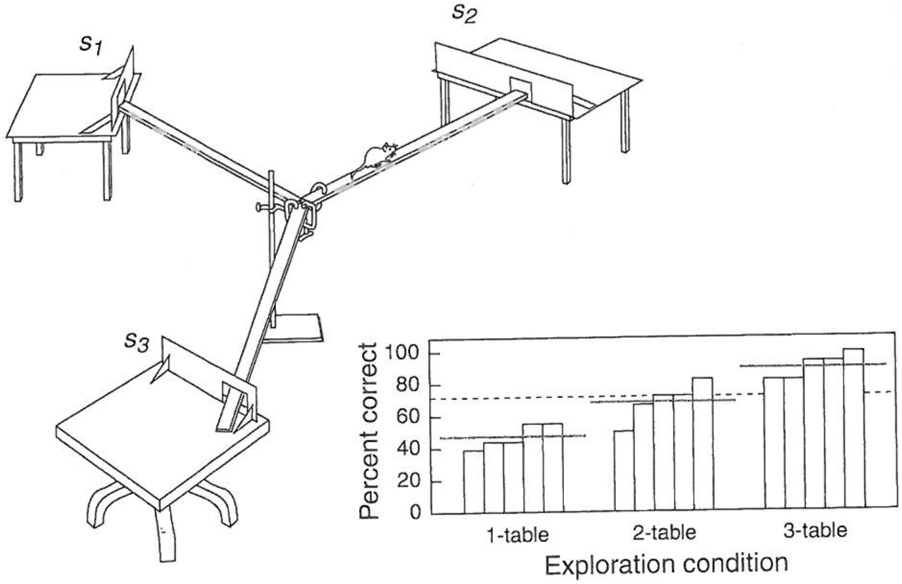

Figure 8.15. Setup for the Maier three-table task, redrawn from Maier (1932b) with permission. Screens are placed so that rats cannot see from the runways whether or not food is on any of the tables. Data redrawn with permission from the experiment of Ellen, Soteres, and Wages (1984) in which rats explored the maze piecemeal, one, two, or three tables at a time before being tested. Results are shown from five rats in each group; solid lines are group means; dotted line is performance level that could occur only 5% of the time or less.

ale's territory earlier
planning their route
Male meadow voles
ntered in a T-maze

have to experience
r what? Some of the
m studies with the
al working memory
in a familiar space
connected by a Y-
, a rat is allowed to
laced on the day's
a few minutes but
r two tables, from
ly once a day, with
ally do quite well,
ut accurate choice
aze (Maier 1932a;

e they have never
) restricted experi-
day to explore the

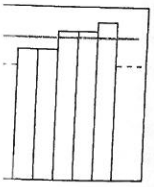

ermission. Screens
the tables. Data
hich rats explored
shown from five
ld occur only 5%

maze, one group exploring only one runway and table per day, one exploring two different connected runways and tables, and one exploring the whole maze. Every three days, the rats were given the standard three-table test. The rats that explored only one runway at a time never performed above chance in 18 such tests, whereas rats given full exploration performed above chance from the outset (Figure 8.15). Thus the information gained from piecemeal exploration does not seem to be knitted together into a unitary representation. Other findings from mazes and swimming pools (e.g., Maier and Schneirla 1935/1964; Sutherland et al. 1987; Save et al. 1996) agree that to treat different places as connected a rat has to travel between them. Seeing they are connected is not enough. Perhaps this finding should not be surprising for an animal that normally does most of its traveling in the dark.

8.3.2 Learning about redundant cues: Competition or parallel processing?

O'Keefe and Nadel (1978) suggested that exploring novel items in a familiar space allows an animal to update its cognitive map in the same way as a cartographer adds a new farmhouse or removes a hedge from a printed map. Incorporating all available cues into a cognitive map would ensure redundancy when primary cues fail, which could be important for tasks like getting home. Indeed, an example of backup mechanisms is illustrated in Figure 8.13. As another example, experienced homing pigeons tested on sunny days use a sun compass, but birds tested under thick cloud cover can home just as well, relying on landmark memory, olfaction, magnetic information and/or infrasound (Keeton 1974). But several of O'Keefe and Nadel's ideas are contradicted by more recent findings in associative learning. One apparent contradiction is the phenomenon of *latent inhibition,* that is, mere exposure to a potential CS may retard later learning about it (Chapters 4 and 5). This is the opposite of what would be observed if the animal is continuously building a cognitive map. However, exposure to a particular spatial context does sometimes retard later learning about locations within it. Just as in associative learning, preexposure enhances discrimination (i.e., perceptual learning occurs) when the locations to be learned about are similar, but latent inhibition occurs when they are very different (Rodrigo et al. 1994; Chamizo 2003). Incorporating redundant cues into a cognitive map is also at odds with the principle of cue competition in associative learning. The Rescorla-Wagner model formalizes this principle, most clearly evidenced in the phenomena of overshadowing and blocking (Chapter 4). So do overshadowing and blocking occur in spatial learning?

Beacons and landmarks

One clear example of cue competition in spatial learning comes from a study of blocking with rats in a water tank by A. Roberts and Pearce (1999). The time spent in the quadrant of the tank where the platform was usually found was compared in two groups of rats (Figure 8.16). Both groups had been trained with a beacon attached to the platform and landmarks visible around the tank, but the blocking group were first trained with curtains drawn around the tank so they learned about the beacon alone. This initial training was expected to block learning about the added landmarks, and Figure 8.16 shows that indeed it did. Further controls with various kinds of swimming experience in the first stage still learned more about the landmarks in stage 2 than rats already trained to use the beacon.

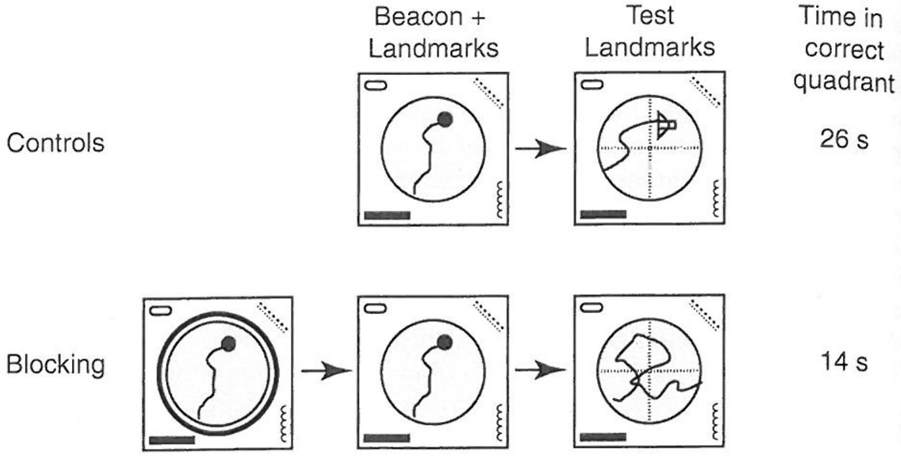

Figure 8.16. Procedure and results of A. Roberts and Pearce's (1999) test of blocking in the water tank. The heavy line around the outside of the tank represents a curtain; the black dot is the dry platform. The final test lasted 60 seconds; hence the blocking group's test score represents random search, i.e. about 1/4 of the time in the quarter of the tank with the platform.

As might be expected, landmarks closest to a goal overshadow more distant landmarks, for example in bees (Cheng et al. 1987), pigeons (Spetch 1995), and rats (Morris 1981; Redhead et al. 1997; Chamizo 2003). Landmarks can also block each other. For example, learning to use a set of three landmarks to locate the hidden platform in a swimming pool blocks rats' learning about a fourth landmark added later on (Rodrigo et al. 1997). Blocking and overshadowing have also been found between intramaze cues (floor texture) and extramaze cues (i.e., landmarks in the room) in a radial maze (Diez-Chamizo, Sterio, and Mackintosh 1985; March, Chamizo, and Mackintosh 1992). Of course such findings are not necessarily incompatible with observations of animals using normally redundant cues in a hierarchical manner. Overshadowing and blocking do not need to be complete. If some cues are simply learned more strongly than others, one would expect those learned best to be used first when available. In any case, landmarks and beacons tap only a subset of spatial processing modules. Tests of overshadowing and blocking may have different results when different spatial modules are brought into play.

Geometry and landmarks

As we saw in Section 8.1.5, when animals are disoriented relative to the outside world, they initially rely on information about the location of a goal relative to the shape of an enclosure and ignore more informative features. Some discussions of the geometric module have suggested that geometry is used for reorientation, perhaps supporting learning about features but not competing with it (e.g., Cheng 1986; Wang and Spelke 2002). That notion suggests that geometric cues should not be overshadowed or blocked by other cues. One test of this idea is illustrated in Figure 8.17. Rats were trained in a rectangular enclosure with a sawdust-filled bowl in each corner and a reward buried in the bowl near a black landmark. Rats learned to go directly to the bowl near the landmark, but they learned about geometry at the same time, as the vast majority of errors during acquisition consisted of digging in the diagonally opposite, that is, geometrically equivalent, bowl. In addition, like

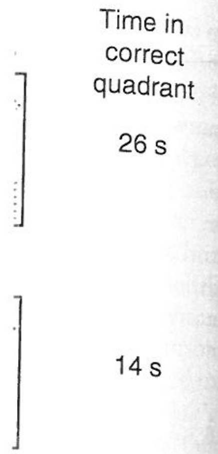

locking in the water
lack dot is the dry
re represents random

more distant land-
h 1995), and rats
an also block each
locate the hidden
h landmark added
e also been found
landmarks in the
sh 1985; March,
necessarily incom-
s in a hierarchical
. If some cues are
learned best to be
p only a subset of
nay have different

ve to the outside
oal relative to the
discussions of the
entation, perhaps
g., Cheng 1986;
es should not be
is illustrated in
a sawdust-filled
k landmark. Rats
d about geometry
nsisted of digging
In addition, like

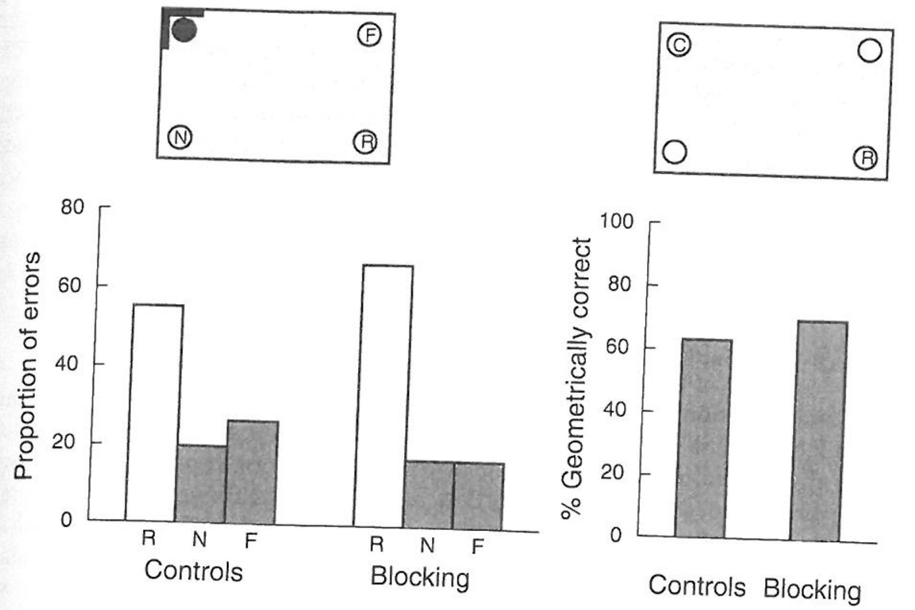

Figure 8.17 Training with a landmark (the black corner) does not block geometry learning. Rats in the blocking group were pretrained with the black corner panel in a square enclosure, yet they still made a high proportion of geometric (rotational, R) errors when transferred to a rectangle (left panel; errors summed over all of training). In a test without the landmark at the end of training, they chose the geometrically correct corners C and R, as often as controls not pretrained with the landmark (right panel). Data from Experiment 3 in Wall et al. (2004).

other species tested in a similar way (Cheng and Newcombe 2005), the rats preferred the geometrically correct corners in a test without the landmark at the end of acquisition, showing they had learned about the geometry as well as the more informative landmark. Moreover, learning based on the shape of the rectangle was not blocked by prior training with the landmark in a square (i.e., geometrically uninformative) enclosure (Wall et al. 2004).

Tests of overshadowing and blocking in water tanks of various shapes have also revealed little evidence of cue competition. Indeed, sometimes a cue at a goal facilitates learning of geometry (e.g., Pearce et al. 2001; N. Y. Miller and Shettleworth 2007). But although these findings suggest that learning the location of a goal relative to the geometry of an enclosure goes on independently of learning about its location relative to features within the enclosure, a deeper analysis reveals that cue competition is still at work (N. Y. Miller and Shettleworth 2007, 2008). Searching for a goal is an instrumental task, so the animal's choices determine the frequency with which the cues at each location searched are paired with reward or nonreward. For instance, when rats begin to learn that the salient black landmark signals reward in the situation in Figure 8.17, they start choosing the marked corner more often. When they do choose that corner they also experience a pairing of its geometry with reward. The learning based on these pairings is reflected both in the relatively large proportion of geometric errors early in training and in geometrically correct choices during tests without the landmark. The same process leads the blocking group to learn about the geometry of the rectangle. Geometry and features are competing for learning, but this is typically not evident in choices because if a location is chosen on the basis of the total associative strength of its cues relative to the total at all locations, one location can be quite strongly preferred over the others even while none of the individual cues is at asymptote. For example, the corner by a landmark in a square enclosure can be chosen a large proportion of the time even if its associative strength is not high enough to block geometric cues when transferred to the rectangle. One prediction of this model, then, is that conventional cue competition will reveal itself in choices under some conditions, including after prolonged training in either stage of an experiment like that in Figure 8.17. The shape of the enclosure and the distribution of features within it will also influence the results, as indeed it does (Miller and Shettleworth 2007).

Dead reckoning and beacons

The characterization of dead reckoning as an obligatory process, a basis for learning the locations of stable allocentric cues yet always going on in the background implies that the idiothetic cues for dead reckoning do not compete with learning landmarks and/or beacons. Indeed, if they did how would animals ever learn about stable allocentric cues? But although some of the findings described in Section 8.2.3 imply that dead reckoning operates independently of beacon and landmark learning, only one study (Shettleworth and Sutton 2005) has tested this implication with a study of overshadowing and blocking. Rats found food pellets in a large circular arena and carried them back to eat in a home cage that was concealed behind one of 16 identical doors on the periphery of the arena. In some conditions the correct door was surrounded by a black panel, functioning as a beacon, and in other conditions no exteroceptive cue identified it so the rats had to home by dead reckoning. To ensure that the rats could not locate the home on the basis of cues outside the arena, the home cage and black panel were in varying locations in absolute space, and the rats were disoriented before entering the arena. Rats trained with the beacon homed no more accurately than rats trained with a "beacon" at random locations relative to the home door, and the two groups homed equally accurately in tests without the beacon. That is, beacon learning did not overshadow the cues used in dead reckoning. And as might be expected on functional grounds, rats that already had extensive experience homing on the basis of dead reckoning alone learned as much about the beacon when it was added as rats that had it from the outset, that is, dead reckoning did not block beacon learning.

Places and responses

Extensive literature also points to a noncompetitive interaction between spatial learning systems in the case of "place" (or landmark) and response learning (sometimes referred to in this context as *habit* learning). Figure 8.18 shows a classic setup for testing whether rats that learn to make a particular turning response to arrive at a particular place in a T-maze have learned the response or the place. Notice, however, that the test consists of forcing the rat to choose between making the rewarded response and going to the rewarded place. A choice consistent with one kind of learning does not mean that the other kind has not also occurred. The typical finding in such tests is that early in acquisition place learning is evident, but later the habitual response prevails (Restle 1957; Packard and McGaugh 1996). This finding in itself suggests cue competition is not going on. If it was, how could response learning develop when place learning was already allowing the animal to locate the reward? Neurobiological data also lend support to the idea that rather than competing, place

ne of the individual cues
square enclosure can be
ength is not high enough
. One prediction of this
l itself in choices under
stage of an experiment
nd the distribution of
ed it does (Miller and

cess, a basis for learn-
on in the background
compete with learning
imals ever learn about
cribed in Section 8.2.3
n and landmark learn-
d this implication with
llets in a large circular
concealed behind one
conditions the correct
beacon, and in other
ad to home by dead
e on the basis of cues
varying locations in
he arena. Rats trained
d with a "beacon" at
roups homed equally
g did not overshadow
n functional grounds,
sis of dead reckoning
rats that had it from
ing.

ion between spatial
onse learning (some-
shows a classic setup
esponse to arrive at a
ce. Notice, however,
aking the rewarded
nt with one kind of
. The typical finding
ut later the habitual
This finding in itself
d response learning
locate the reward?
an competing, place

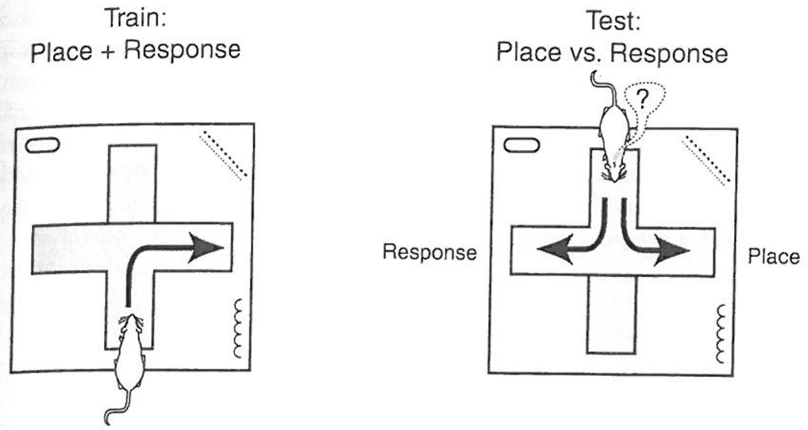

Figure 8.18. Test of whether rats trained in a T maze as indicated in the left panel learn where to find reward or what response to make. The two possibilities are dissociated as indicated in the setup on the right. The dark bar is a block that confines the rat to one part of the maze. The random shapes around the edge are objects in the room that can serve as landmarks.

and response learning go on in parallel, in different brain regions, the hippocampus and caudate nucleus respectively (White and McDonald 2002).

This was beautifully demonstrated in a study using a setup like that depicted in Figure 8.18 in which rats were given the place versus response test early or late in place + response training and while the hippocampus or caudate was temporarily disabled with injections of lidocaine (Packard and McGaugh 1996). Rats with the hippocampal place-learning system disabled chose randomly in the early test but chose on the basis of the trained response in the later test, consistent with the observation that response learning emerges gradually. In contrast, rats with the caudate habit system disabled chose the rewarded place in both the early and late tests. Thus response learning did not emerge at the expense of place learning: it remained intact and could be revealed when the competing behavioral tendency was removed. The same issue was addressed without pharmacological manipulations by testing rats in a radial maze arranged so that response learning and place learning could be dissociated (Gibson and Shettleworth 2005). The results suggested that prior response learning interferes with learning about landmark (place) cues introduced later. This may have occurred because rats for which a habitual response continues to lead to reward when place cues are added simply pay less attention to, or spend less time exposed to, those cues than controls.

Redundant cues in spatial learning: Conclusions

The picture of spatial learning sketched here is not that of a single system in which cues compete for a limited amount of predictive value like CSs in conditioning. Landmarks compete for learning with beacons and geometric cues and with each other, but dead reckoning goes on in parallel yet somehow in support of learning about allocentric cues. At least in mammals, response or habit learning is another parallel system. The existence of parallel systems allows for the redundancy which functional considerations suggest is particularly important for spatial tasks like getting home. Ultimately whether different mechanisms evolve so as to compete would be expected to reflect the relative costs and benefits of relying on minimal sufficient predictors versus processing redundant information. In addition, the costs and benefits of any one mechanism may be balanced by the costs and benefits of others. For instance, well-learned responses demand little attention and permit fast travel in familiar places, but a slower, more attention demanding system such as exploration and landmark learning is called for when conditions change. Evolutionary pressure to optimize over different criteria may account for the widespread existence of multiple spatial learning systems. Whether the multiple kinds of memories that result are integrated into a cognitive map in any sense is the question for the next section.

8.4 Do animals have cognitive maps?

8.4.1 What is a cognitive map?

In Chapter 4, "How does the animal represent the CS?" meant simply, "What features of the CS are encoded or remembered?," a rather minimal kind of representation. In contrast, the representation embodied in a cognitive map is typically assumed to encode distances and directions and to enable mental operations on them. To take an example we will shortly consider in more detail, an animal that can encode the distance and direction of two feeding sites from a home base and whose nervous system is capable of implicit computations analogous to the operations of vector algebra can move directly between the two feeding sites without going home in between (Figure 8.14, "metric map").

Distances and directions are the *metric properties* of space. Blueprints, city plans, road maps, and globes are useful because they represent distances and directions accurately. But plenty of useful maps do not preserve such vector information. A familiar example is a subway route map. Such a map is useful for planning a trip on the subway because it shows which station is on which route and what order they can be reached in. Such a *network map* can be used without its representing distances between stations or angles between connecting routes. Indeed, because these may not be represented accurately, a tourist wanting to explore the city on foot would be foolish to use a it as a guide. In contrast to the subway map, a map that preserves distance and direction information, a *vector* map, allows the planning of novel routes to unseen goals. How useful it is, though, depends on the density of identifiable locations represented. For example, a tourist starting from an obscure side street armed only with a vector map of the city landmarks has to wander around until finding a place marked on the map. This potential limitation of real Euclidean paper maps has traditionally been overlooked. It is an assumed unlimited flexibility that has distinguished a cognitive map from "mere" reliance on one or more kinds of spatial cues. As Section 8.2 shows, however, single cues or combinations of them can guide animals very effectively. This means that it is almost impossible to find indisputable evidence that any animal is using a cognitive map in the sense of a global representation of space equivalent to an overhead view that preserves distances and directions among an infinity of locations. Whether *cognitive map* always means the same thing is a problem too, as we see by surveying some of the landmarks in its history.

Tolman

E.C. Tolman introduced the term *cognitive map* into psychologists' debate about place versus response learning. Rather than simply acquiring chains of stimulus-response connections, Tolman claimed, rats in mazes acquire "something like a field map of the environment" (Tolman 1948, 192). Stimuli influence behavior not through S-R connections, but through the mediation of the cognitive map. Cognitive maps could be broad and comprehensive or "narrow strip maps," confined to knowledge of specific routes. The most compelling data Tolman cited in support of his hypothesis were from tests of *latent learning* and ability to take novel shortcuts in mazes. In a typical latent learning experiment, a rat was allowed to explore a maze without receiving any reward. For instance, food might always be present in one location but the rat would be satiated. If the rat ran straight to the food when it was hungry later on, its behavior could not have resulted from the reinforcement of S-R connections because it had not been getting any reinforcement. Therefore, it must have learned the location of the food and generated appropriate behavior on the basis of this knowledge. Similarly, an animal that took an efficient novel shortcut when displaced to a new location or when its usual path to a goal was blocked must have acquired knowledge about the goal as a place.

Much of the behavior just described is now largely taken for granted in the view of conditioning as a representational process described in Chapter 4, so in a sense Tolman's view of learning as S-S connections has prevailed. It is no longer problematical, for instance, that animals approach or avoid places on the basis of knowledge about their value. Indeed, this is the basis of the popular *conditioned place preference* test. In this paradigm an animal is first exposed to each of two distinctive chambers in each of which a different biologically significant event occurs, for example food in a grey square chamber and opportunity to run in a wheel in a striped round chamber. The relative value it gives to them is then measured by removing the rewards, connecting the two chambers and seeing where the animal spends more time. And in at least some of Tolman's experiments, tests of rats' specifically spatial knowledge did not go beyond demonstrating such S-S learning. For instance, if cues near the goal were still visible from a novel starting point, rats could approach them without any maplike knowledge, as in the experimental arrangement depicted in Figure 8.19 (Tolman, Ritchie, and Kalish 1946).

O'Keefe and Nadel

Although Tolman's views are important in the history of psychology, he actually said rather little about the properties of cognitive maps and how they might be acquired. After a lapse of 30 years, this gap began to be filled by John O'Keefe and Lynn Nadel (1978) in their influential book *The Hippocampus as a Cognitive Map*. They developed the view that some organisms, including humans, rats, and migratory birds, possess cognitive maps, in the sense of a unitary, allocentric, connected spatial representation in which experience locates objects and events. Acquisition and use of the cognitive map is supported by the locale system, a cognitive module located in the hippocampus of vertebrates. The locale system contrasts with the taxon system, which supports conditioning and is located elsewhere in the vertebrate brain. The learning supported by the taxon system was seen by O'Keefe and Nadel as relatively inflexible compared to that supported by the locale system, but developments in the last 30 years have undercut this distinction. As in Tolman's account, maps are acquired through exploration (latent learning) and allow more flexible behavior than route learning. However, as in the example mentioned above, some flexibility can result from stimulus generalization, and at least over relatively short distances, path integration allows

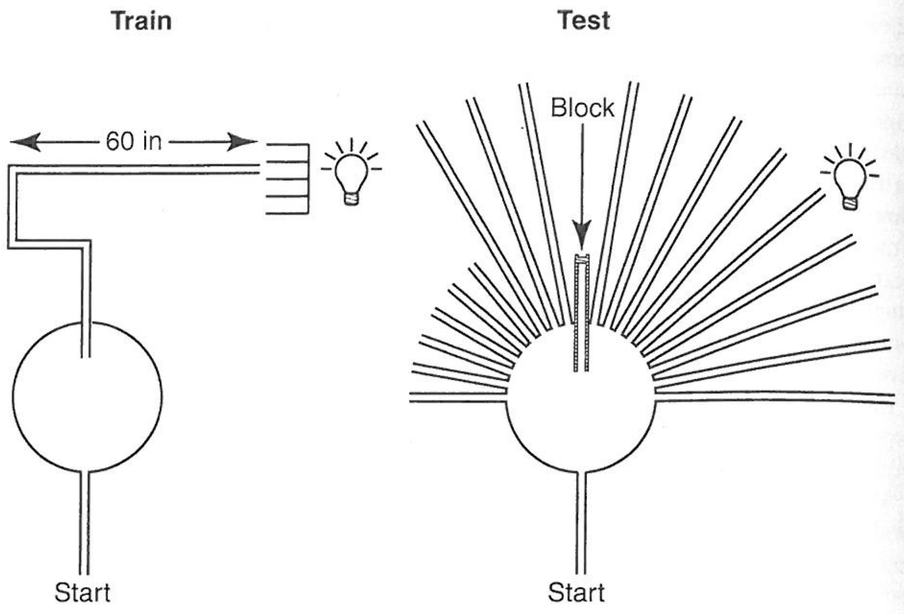

Figure 8.19. Setup used by Tolman, Ritchie, and Kalish (1946) to test place learning in rats. The alleys were arranged as shown on the left for training; for testing they were replaced with the "sunburst" maze shown on the right. Rats tended to choose the new path that led directly toward the goal. Note the light at the goal box. Redrawn from Tolman et al. (1946) with permission.

travel toward unseen goals in a way not fully appreciated when O'Keefe and Nadel first wrote.

Gallistel

Gallistel's (1990) review of spatial behavior in animals has a very different flavor from O'Keefe and Nadel's, partly because by 1990 the cognitive revolution in psychology had made the notion of cognitive mapping more acceptable. In addition, this period witnessed an explosion of relevant research, and Gallistel was the first person writing for people in the cognitive sciences to bring together the new laboratory studies on rats in the Morris swim task and radial maze with biological field work on bees, ants, homing pigeons, and other animals. This integrative approach has been tremendously influential and is now almost taken for granted, as can be seen in numerous more recent reviews (e.g., Newcombe and Huttenlocher 2000; T. Collett 2002; Jeffery 2003). Gallistel's (1990) definition of *cognitive map* is fairly loose (e.g., Chapter 6, 121): any orientation based on implicitly computing distances and directions is evidence of a cognitive map. Dead reckoning, matching "snapshots," or responding to landmarks all count as cognitive mapping, albeit perhaps in a small-scale and limited way. Species may differ in the richness and detail of their cognitive maps, but evidence for them is ubiquitous.

Bennett

Because *cognitive map* means different things to different people and because most frequently used behavioral assays of cognitive mapping have not ruled out

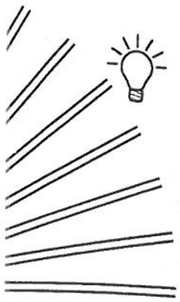

ırning in rats. The
laced with the
ɟ directly toward the
ermission.

ıen O'Keefe and

ⁱ different flavor
ve revolution in
ıble. In addition,
ƨtel was the first
the new labora-
ı biological field
;rative approach
d, as can be seen
2000; T. Collett
fairly loose (e.g.,
ances and direc-
"snapshots," or
haps in a small-
f their cognitive

e and because
ᵎ not ruled out

---

well-defined alternatives such as dead reckoning or generalization from familiar local views, it is almost impossible to find unambiguous evidence for it. Discussion of cognitive maps should be replaced with better-grounded specification of how animals (including people) find their way from place to place. This position was stated forcefully by Bennett (1996) and echoed by Mackintosh (2002) in an article titled, "Do not ask whether they have a cognitive map but how they find their way about." To see why this is good advice we need to review the approaches to testing for cognitive maps.

8.4.2 Mapping and short-range orientation

Shortcutting

A central behavioral prediction from any notion of cognitive mapping is that within familiar terrain an animal with a cognitive map should be able to reach a goal by a novel route. It will take a novel shortcut when one is made available, and if it is displaced to a new starting place it will head directly to the goal rather than returning to the familiar start before continuing its journey. Tests of this prediction have a long history, beginning with the work of Tolman, Ritchie, and Kalish (1946) illustrated in Figure 8.19. In laboratory studies like theirs it is easy to guarantee that the offered shortcut is novel, but it is not so easy to be sure that animals are doing anything other than orienting by landmarks (and indeed, that rats use landmarks rather than only S-R habits may have been all that Tolman aimed to establish). As long as landmarks visible from the goal are visible to the animal in the same left-right relationship at the point where it chooses between the shortcut and some other route, the self-to-goal vector computed as in Figure 8.7 will take the animal along the more direct route. Cues at the goal clearly influenced the rats in Tolman et al.'s (1946) original study, because there was a distinctive light right at the goal box (see also Chapuis, Durup, and Thinus-Blanc 1987). In a careful study in which dogs often took the shortest route between two novel locations in a large field (Chapuis and Varlet 1987), this could have been a problem too. The dogs were led first to one location and then another and shown meat in each one before being released from the common starting place for these trips to find the food. As they were shown each piece of food they could have encoded its location with respect to features in the surrounding familiar environment. Indeed, rats can acquire new knowledge about what is where in a single trial in an environment which already supports a network of associations (Tse et al. 2007).

These considerations mean that to test whether animals are using a representation that includes more than local landmark-goal vectors, landmarks perceptible from the goal must not be perceptible when the shortcut is chosen. The importance of this requirement is very well illustrated by a much-discussed series of studies with honeybees. Honeybees are ideal subjects for studies of spatial orientation in natural landscapes because foragers routinely make many round trips each day between the hive and feeding sites hundreds of meters away. Using methods pioneered by Karl von Frisch (1967), marked individuals can be trained to artificial feeding sites selected by the experimenter. Newly emerged foragers gradually become familiar with the area around the hive, as shown by the observation that when bees are released some distance from the hive, the experienced individuals are more likely to find their way back (review in Dyer 1994). Extensive and detailed spatial knowledge thus seems to exist in bees' tiny brains. Discussion of whether it can be described as a cognitive map has centered around a shortcut experiment originally reported by Gould (1986) and repeated by others, sometimes with different results (Wehner and Menzel 1990; F. Dyer 1991, 1994). Bees were trained to only one of two feeders, F1 and F2, equidistant from the hive but out of direct sight of each other. The lines connecting A, B, and the hive formed an approximately equilateral triangle, as in the arrangement depicted in Figure 8.14. The test of whether the bees knew the relationship of a feeding site to the landscape as a whole consisted of capturing marked individuals as they left the hive for one site, say F2, and releasing them at the other. A bee released at a novel location flies up maybe 9 or 10 meters, circling around as if getting its bearings, and then heads off in a definite direction. Data in these studies thus consisted of the compass bearing recorded for each bee when it vanished from view.

Because when tested Gould's bees tended to head toward the site they had been trained to, he concluded that the bees had a "maplike representation" of their local environment. But because their experience was not controlled, it is impossible to know whether the shortcut was truly novel. Moreover, when the bees flew up to get their bearings they could have gotten a view of the landscape sufficiently similar to that near the familiar feeding site to allow them to orient. And indeed, "maplike" orientation in such a test does seem to require this, as shown by Dyer (1991) in a similar study that had one important difference. One of the two critical feeding sites, B, was down in a quarry whereas A was up at the same elevation as the hive. Bees trained to B, in the quarry, and released at A behaved like Gould's bees and headed off from the novel release site toward the feeding site B. But bees trained to feed at A, on the high ground, and released at B, in the quarry, could not easily get a view similar to that which they saw when leaving the hive for A. These bees did not head either for the hive or for A but departed from B in the same compass direction they had been taking when they left the hive. This did not reflect some peculiarity of site B in the quarry; bees trained to fly directly between A and B were able to orient accurately.

As we see presently, this is not the end of the story of cognitive mapping in bees, but Dyer's study remains an important demonstration of why tests of shortcutting must ensure that cues at the goal are not perceptible at the choice point. This was done for rats in the enclosed maze depicted in Figure 8.20 (Singer, Abroms, and Zentall 2006). The three goal boxes were identical, but each arm had a distinctive

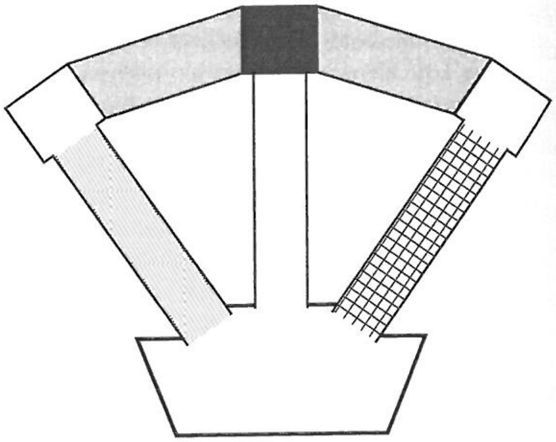

Figure 8.20. The enclosed maze used by Singer, Abroms, and Zentall (2006) to test cognitive mapping in rats. Different shadings signify different textures in the maze arms. The arms linking the three goal boxes (top of the figure) were open only for testing. Redrawn from Singer, Abroms, and Zentall (2006) with permission.

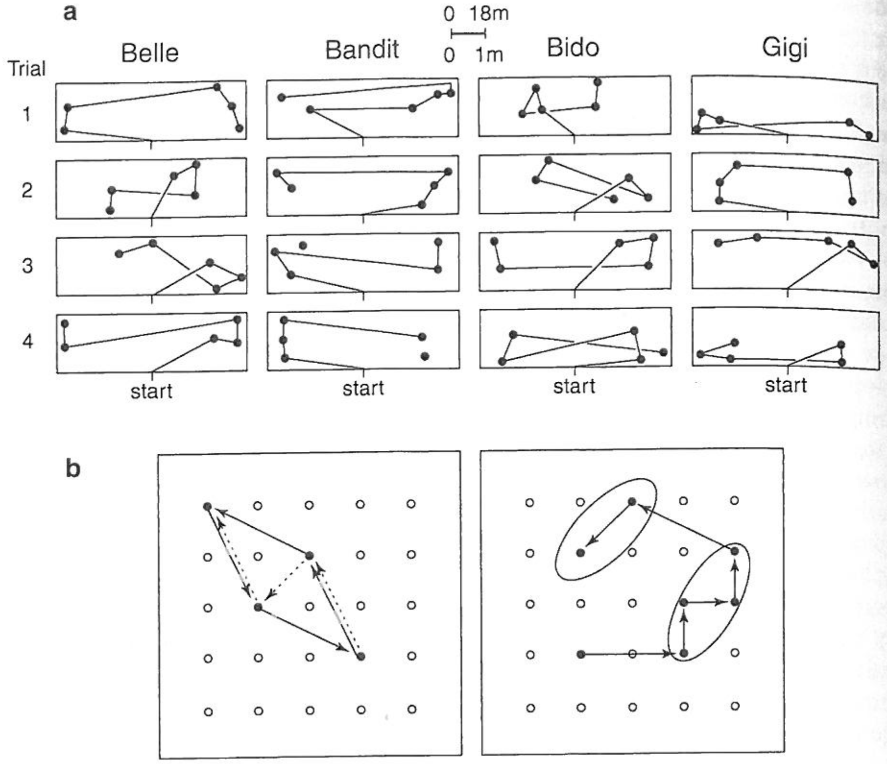

Figure 8.21. a. Paths taken by each of four chimpanzees in their first four trials with five hidden food items. Redrawn from E. Menzel (1978) with permission. The marker indicates the scale in both a and b. b. "Diamond" and "unequal sides" configurations used for testing spatial planning by vervet monkeys. In the diamond configuration, the shortest path among all four corners is the dotted one if the animal is not planning to return to the start (lower right vertex) but the black route if it is. Redrawn from Gallistel and Cramer (1996) with permission.

Nevertheless, cognitive mapping (or at least planning a route beyond the next move) can in principle be distinguished from wholly local choices with certain arrangements of sites. For example, the optimal path for collecting food from the diamond-shaped arrangement shown in Figure 8.21b depends on whether the animal is going to return to the starting point. The animal can make the optimal choice after the second food item only by planning beyond the next two choices. Similarly, when four food items are on one side of an arena and two on the other, the animal must mentally look beyond the first two items to be collected in order to move optimally. Cramer and Gallistel (, 1997; Gallistel and Cramer 1996) report that vervet monkeys behaved as if planning routes in both of these tests. However, without knowing the animals' reinforcement histories in the testing situation, which are not reported, it is difficult to know how to evaluate these data. Moreover, in the four- versus two-item test, the monkeys might simply have remembered the area with four items better if they spent more time there while the sites were being baited. Nevertheless, this approach could be pursued further than it has been (see also Janson 2000).

Some species of jumping spiders are well known for their ability to choose efficient routes and detours in natural conditions. These spiders do not weave a

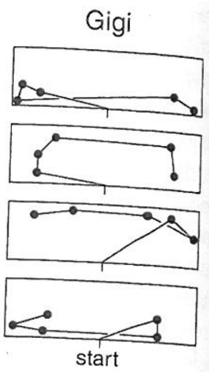

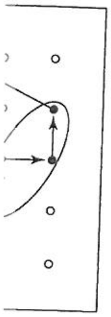

als with five hidden food
es the scale in both a and
l planning by vervet
rners is the dotted one if
black route if it is.

ite beyond the next
hoices with certain
cting food from the
whether the animal
optimal choice after
ces. Similarly, when
er, the animal must
to move optimally.
hat vervet monkeys
ithout knowing the
e not reported, it is
ur- versus two-item
four items better if
Nevertheless, this
on 2000).
ability to choose
's do not weave a

---

web but pursue their insect prey visually, sometimes watching a victim for a long time while slowly creeping up on it. The spider might move away from the prey initially and climb up a branch from which it can pounce. Tarsitano and Andrew (1999) captured this situation in the laboratory by placing a spider (*Portia labiata*) in the middle of a square enclosure where it could see a prey item above and in front of it. To reach the prey, the spider had to walk diagonally away from it, climb a pole, and traverse a series of two ramps. When spiders were confronted with a choice between two such routes, one to each side, and one of them had a gap in it, 16 of 18 spiders headed in the direction of the unbroken path as soon as they left the start platform. When both routes were complete, they chose the one they had scanned the most while sitting on the start platform; when one route was incomplete, simple algorithms describing scanning ensured they spent most time scanning the complete one. Like the ants we meet shortly, the spiders provide a nice illustration of how apparently demanding feats of navigation can be accomplished by simple mechanisms (see also Cheng 2006).

Knitting together

The idea that animals orient with reference to a cognitive map implies that information gathered in different parts of a journey, even qualitatively different kinds of information, is knitted together into a single allocentric representation. Unlike in the examples of integrating simultaneous cues in Section 8.2, here cues are encountered successively. This occurs in a setup developed by Benhamou (1996) for rats in a water tank and subsequently adopted by Gibson and Kamil (2001; Gibson 2001) for Clark's nutcrackers and people (Figure 8.22). The subject's task is to find a goal hidden at a fixed location in a room full of landmarks. The subject encounters the landmarks on the first part of the

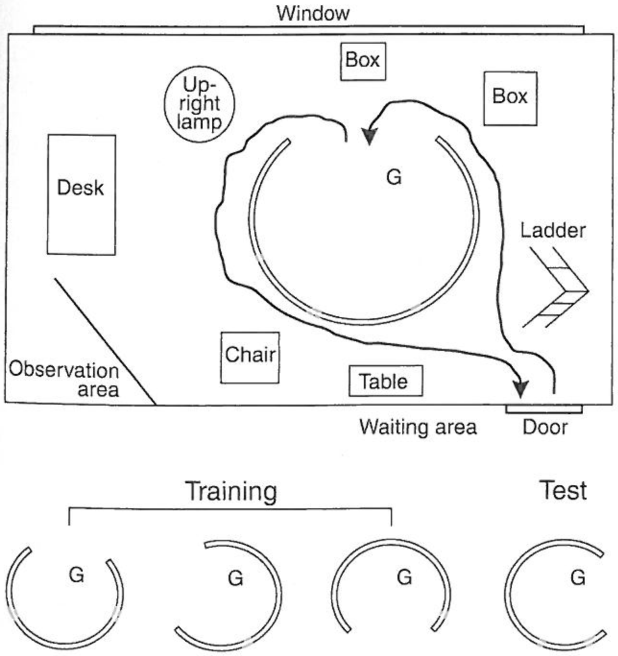

Figure 8.22. Overhead view of the setup used by B. Gibson (2001) to test cognitive mapping in people. The semicircle indicates the opaque enclosure within which the goal (G) was located. The enclosure was rotated from trial to trial to allow about a 270° view of the room from the goal in total. Three of the positions used are shown below the diagram, along with the orientation used for the test. The arrows are the paths people used to enter and leave the enclosure. Redrawn from B. M. Gibson (2001) with permission.

journey, but the goal is within an enclosure that permits only about a 90° view of the room. This enclosure is rotated from trial to trial, so that its entrance has no consistent relationship to the goal, and so that over trials the subject sees different parts of the landmark panorama from the goal. Once subjects can find the goal they are tested with the enclosure in a novel orientation, so they have a new view of landmarks from within it. Neither rats nor nutcrackers oriented accurately in initial tests of this sort, though some nutcrackers eventually learned to do so. Moreover, people behaved similarly if, like nutcrackers, they had to learn for themselves that the goal was at a constant location in the room. Rather than developing a representation of the goal within the room, subjects of all species tested relied much more than they should have on local cues such as the vector between the edge of the enclosure and the goal. In larger-scale space too, people do not do a good job of relating their orientation within an enclosed area such as a room to that within the surrounding environment (Wang and Brockmole 2003).

Some different results have come from another test of knitting together devised by Blaisdell and Cook (2004) for pigeons. As in second-order conditioning in learning of temporal relations (Section 4.4.3), the animal is first exposed to a relationship between two neutral stimuli, simultaneously presented landmarks rather than successively presented tones, lights, or the like. Then it learns to locate food with respect to one of those landmarks, and finally, it is tested with the other. For example, suppose in the first phase A is west of B, and in the second, food is south of A. Knitting together these two experiences would lead the animal to search southwest of B in the test. Generalizing from A to B, a possible alternative strategy, would be expressed as searching directly south of B. Pigeons behave as if connecting the two experiences, whether the landmarks are presented in an open field (Blaisdell and Cook 2004) or on a touchscreen (Sawa, Leising, and Blaisdell 2005). However, in a video simulation of the open field task people at first show generalization between the A and B landmarks (as indeed the pigeons did in the actual open field), but gradually transfer their searching to the site specified by integration (Sturz, Bodily, and Katz 2006). This pattern may reflect learning from initial nonreinforced searches in the test rather than a mapping-like process. However, rats tested in a slightly different way do seem to knit together separately experienced items of landmark information (Chamizo, Rodrigo, and Mackintosh 2006). Rats were trained to find the hidden platform in a water tank in intermixed trials with two sets of three landmarks having one member in common (e.g., landmarks A, B, C and C, D, E). Unlike control rats trained similarly but with nonoverlapping sets of landmarks, those for which the sets shared a member preferred the part of the tank with the platform when tested with a novel combination of landmarks (e.g., A, B, E). The variety of results here indicates that much remains to be done to understand the extent to which animals knit together separately experienced spatial relationships into an overall "map."

Australian desert ants (*Melophorus bagoti*), however, fail entirely and in a surprising way to knit together information obtained in different parts of a journey. The ants are an instructive corrective to anthropomorphism, a reminder of how almost unimaginably strange ways of navigating can be perfectly effective in the natural conditions in which they are normally used. Unlike the Tunisian desert ants we met in Section 8.1, the Australian species live among grass tussocks that provide landmarks, which the ants evidently use along with global path integration somewhat as illustrated in Figure 8.12. By means of a system of barriers, Wehner and colleagues

(Wehner et al. 2006) forced ants to adopt different outward and homeward paths across such terrain to a constant food source. Within a few trips, each ant developed an idiosyncratic round trip route. Having thus shown that they knew the way to the food and home again, the ants were picked up while on the way home with a biscuit crumb and placed down partway along the outward path. What they did then depended on how close to home they were when picked up, but in no case did ants behave as if recognizing where they were on the outward path by either heading back along it or taking a shortcut home. Rather, ants that were still some way from home ran in the direction of the global vector that would have led them home from where they were collected. Ants whose global homing vector was at zero because they were caught just before entering the nest behaved as if lost, searching in circles until they hit the homeward route. To quote the authors' summary, "familiar landmarks are not decoupled from the context within which they have been acquired and are not knitted together in a more general and potentially map-like way. They instruct the ants what to do rather than provide them with map-like information about their position in space" (Wehner. et al. 2006, 75).

Cognitive maps in bees revisited

Dyer's bees tested in the quarry behaved similarly to the ants deprived of familiar landmarks by orienting along the vector that would have taken them to their destination from their original starting point. Indeed, for bees commuting between the hive and a customary feeding site, running off a fixed vector back and forth is very efficient. But what happens when vector information tells displaced bees they should already have arrived? It turns out that at about this point, some maplike knowledge takes over. Menzel and his colleagues (2005) captured experienced bees as they were about to start home from a feeder and quickly fitted them with harmonic radar antennas before releasing them within about 500 meters of the hive but at a different direction from it. The records of displaced bees' entire homeward paths so obtained show that bees trained with a feeder in a stable location flew directly away from the release site in the direction they would have taken to return to the hive (Figure 8.23). Bees that had attended the dance (see Chapter 14) of a bee returning from the stable feeder behaved similarly. After perhaps a few hundred meters this straight flight gave way to a circuitous searching flight, which was then followed by a second phase of straight flight, headed directly to the hive or the feeder. Circling appeared to allow the bees to recognize some features of the landscape from which they knew the homeward vector. Bees that had been trained to a feeder at varying locations immediately began searching flights.

As Cheng (2006) points out, these findings need not mean that bees have an exhaustive knowledge of places around the hive because a bee finding itself with a view similar to that from two familiar locations would presumably still wind up at its goal by generalizing and averaging the resulting vectors. Nevertheless, as R. Menzel et al. (2005) conclude, the bees' behavior implies that they have maplike knowledge in the form of learned vectors from a variety of familiar locations to the hive and/or the feeder, that is, a vector map. That is to say, the bees have evidently learned vectors linking certain known locations and can compute routes home from them, but they do not necessarily have a comprehensive metric map of their territory. Moreover, they apparently do not reverse the process, in that when told a vector along which to head from the hive in the form of another bee's dance, they do not behave as if

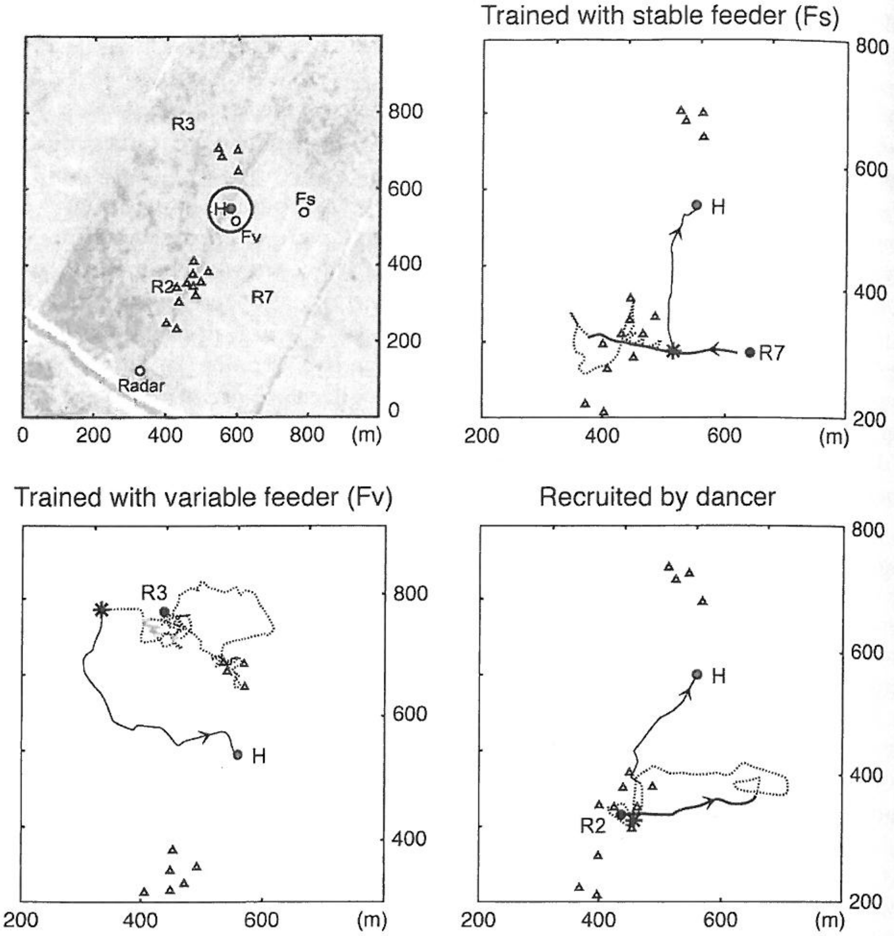

Figure 8.23. Top left: overhead view of the landscape in which bees were tested for maplike knowledge, showing location of the hive (H), the feeder for the stable feeder group (F_S), the tent landmarks (triangles), and several of the experimental release points (R). Remaining panels show paths of one bee from each of the three training groups. Dark lines signify a straight initial flight from the release point. After R. Menzel et al. (2005) with permission.

imagining the location danced about but rather fly to that location even if it is an implausible place for food (see Chapter 14; Wray et al. 2008).

8.4.3 Vertebrates mapping their home ranges

As we have seen, cognitive mapping was originally tested with rats in the laboratory with little or no reference to what the animals might be doing in nature. Indeed, because wild rats are nocturnal and tend to travel along habitual routes and paths (Chitty and Southern 1954), it is not clear what role the visual orientation commonly tested in laboratory rats might have in nature. Ants and bees are more appropriate subjects because their experience can be manipulated in the field and their behavior observed on spatial scales representative of normal foraging trips. As yet no vertebrate has lent itself to such a rich body of work, but a sample of studies of homing pigeons and mammals indicates possibilities for future investigations.

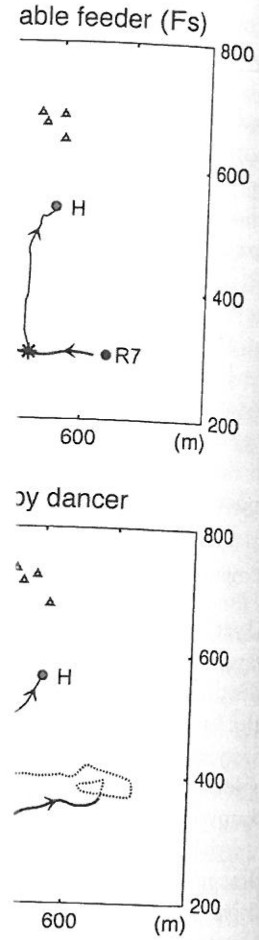

tested for maplike
er group (F_S), the tent
Remaining panels show
straight initial flight from

ation even if it is an

ats in the laboratory
g in nature. Indeed,
ual routes and paths
ientation commonly
re more appropriate
d and their behavior
ips. As yet no verte-
f studies of homing
tions.

Maps and routes in homing pigeons

Given that people have exploited pigeons' homing ability for thousands of years (Wiltschko and Wiltschko 2003), how pigeons find their way home is still surprisingly controversial (Wallraff 2005). The prevailing view is that they use a "map and compass" mechanism, that is, a way to recognize where they are combined with knowledge of which compass direction to head to reach the home loft. Experiments with clock-shifted birds have established that the directional information may be provided by the sun compass (Box 8.2), but other information is important too. When pigeons start home from unfamiliar locations far from their loft, their "map" is in effect a sense of position relative to home based on olfactory, magnetic, or possibly auditory cues (see Wallraff 2005). Remarkable though it is, this is not a cognitive map in the usual sense. And when pigeons home repeatedly from the same site 7 to 10 kilometers from home, they may not refer to a maplike representation either. Rather, data from tiny GPS trackers carried by flying pigeons show that they may develop stereotyped, idiosyncratic, routes. Somewhat like the Australian desert ants or honeybees, when they are released off their usual routes, they first head not toward home but toward the habitual homeward route (Biro, Meade, and Guilford 2004; Meade, Biro, and Guilford 2006). Under the conditions of this study in the richly detailed countryside around Oxford, the pigeons' usual route seems to be encoded as a series of landmarks or views of the landscape. However, the distances involved and other factors apparently influence whether pigeons develop such stereotyped routes (Wiltschko, Schiffner, and Siegmund 2007). In addition, evidence that clock-shifted birds follow their sun compass rather than a familiar route (Wiltschko, Siegmund, and Stapput 2005) indicates that much remains to be understood about the way in which pigeons integrate different sources of navigational information. When orienting in a familiar landscape, they likely use many more kinds of cues than can ants or bees, relying on different ones according to the circumstances (Keeton 1974; Wallraff 2005).

Meerkats finding boltholes

Meerkats (*Suricata suricatta*) are a species of social mongoose found in the southern African semi-desert. We meet them again in discussing social learning and communication (Chapters 13 and 14), reflecting the fact that one South African population has been intensively studied for many years (cf. Ross-Gillespie and Griffin 2007). Meerkats are primate-like in that they form stable social groups with overlapping generations living in a more or less permanent territory. Animals with such a social system obviously have ample opportunity acquire detailed knowledge about what is where and how to get there. Scattered through its 2–4 square kilometer territory, each group has two or more burrow systems for sleeping and raising young, but they also have numerous boltholes into which they can run for safety when threatened by a predator. By observing what meerkats did in response to naturally occurring and recorded meerkat alarm calls, Manser and Bell (2004) showed that the animals know the locations of boltholes. For example, they headed for the nearest bolthole 83% of the time, whether or not it happened to be one they had recently passed while foraging. Meerkats ignored new, human-made, boltholes, even when these were closest to them when an alarm call was heard. In constrast, if the nearest bolthole was one of their own that the experimenters had covered over with a car mat and sand, the meerkats ran to it and tried to get in. Thus their orientation is based on memory, not visual or olfactory cues from an open bolthole. Since a group may have hundreds of boltholes, these findings suggest that the meerkats have extensive and detailed spatial knowledge of their locations. Perhaps this is not surprising, given that they spend many hours each day moving through the territory digging for invertebrate prey and watching for predators. Exactly how they know the positions of nearby boltholes remains to be determined. Columbian ground squirrels (*Spermophilus columbianus),* another small burrowing mammal of comparatively open spaces, locate escape burrows primarily using global cues from the distal panorama of trees, mountaintops, and the like (Vlasek 2006).

In a further test of meerkats' spatial knowledge, on six occasions Manser (personal communication) captured an adult meerkat foraging with its group and released it about a kilometer away at one of the sleeping burrows within the group's territory. The animals traveled quite directly back to the location where they had been captured, arriving within 40 minutes on average. If the group had already moved on, the lone meerkat began looking around, sniffing the ground, and engaging in similar behaviors indicative of searching. The animals were not obviously using vocal or olfactory cues to find the place where they had been removed from the group, but apparently relying on visual recognition of local or global aspects of the landscape.

Monkeys mapping their home ranges

Like the meerkats, apes and monkeys seem likely to possess large-scale integrated representations of space. Field studies of monkeys and apes that involve following habituated groups on a daily basis have provided an abundance of information about the spatial and temporal distributions of the animals' food and how the animals travel between sleeping sites, water holes, fruiting trees, and other resources (Boinski and Garber 2000; Noser and Byrne 2007). The cognitive demands of tracking temporary and spatially dispersed food sources have been proposed to explain differences in relative brain size among primates, with the fruit-eaters supposedly needing larger brains than leaf-eaters (Chapter 12). Numerous field experiments have shown that various monkey species can learn the locations of artificially provided foods and travel among them in an efficient way (e.g., Janson 1998; Garber 2000). One clever study suggests that Japanese macaques (*Macaca fuscata*) remember the locations of trees with a favorite fruit from year to year. C. Menzel (1991) placed akebi fruit, chocolate, or nothing beside a troop's foraging route at a time of year when akebi fruit were not naturally ripe. Monkeys that discovered akebi often left the troop's foraging route and began looking up into akebi trees, whereas those that found chocolate searched the ground nearby.

Observations of unmanipulated animals are necessary for showing how such learning influences their daily travels. These are often very suggestive but must be interpreted with care (Janson 2000; Janson and Byrne 2007). Mapping where a troop of monkeys goes between leaving its sleeping site in the morning and returning in the evening may suggest the animals are planning their routes, but just as with examples of travel among multiple sites in a smaller space, planning ahead has to be distinguished from moving to the next nearest resource on the basis of locally perceptible cues. For example, to decide whether memory for a fruiting tree's location is being used, it is necessary to know the distance from which it can be detected directly, and that differs for forest and savannah species and with the thickness of vegetation across

ientation is based on
nce a group may have
ts have extensive and
surprising, given that
y digging for inverte-
now the positions of
n ground squirrels
nal of comparatively
cues from the distal

x occasions Manser
ging with its group
; burrows within the
the location where
ge. If the group had
sniffing the ground,
ne animals were not
here they had been
ognition of local or

rge-scale integrated
at involve following
f information about
w the animals travel
ources (Boinski and
tracking temporary
plain differences in
sedly needing larger
ts have shown that
provided foods and
r 2000). One clever
ber the locations of
placed akebi fruit,
of year when akebi
ften left the troop's
s those that found

showing how such
gestive but must be
oping where a troop
and returning in the
st as with examples
d has to be distin-
locally perceptible
's location is being
tected directly, and
of vegetation across

---

the year. The animals' nutritional needs relative to different available foods may help explain the sequence in which sites with those foods are visited. Constraints on travel such as needing to arrive at a safe sleeping site by nightfall may need to be taken into account. Noser and Byrne (2007) provide one example of attempting to deal with all these factors. Observed routes can be compared to those predicted on various models of random search (see Janson and Byrne 2007). And ideally, opportunistic observations can help to reveal how the animals' travels reflect what they know. For example, if a predator is encountered at a habitual waterhole, does the troop use another one the next day? In a similar way, the nature and rate of change in travel patterns with the seasons may indicate whether the animals are planning routes with certain goals in mind. For instance, visiting patterns and speed of approach indicate that mangabeys keep track of whether trees are already finished fruiting or are about to produce ripe fruit (Janmaat, Byrne, and Zuberbühler 2006). Such questions have been attracting increasing interest at the same time as they are being addressed in more sophisticated ways. One possibility (Janson and Byrne 2007) is that some primates' spatial knowledge encodes important locations and the routes or vectors between them, but without being a complete Euclidean survey map. Although it would likely represent more and varied types of sites than the vector map of the honeybees, it could be similar in kind.

8.4.4 But do people have cognitive maps?

Research on spatial cognition in human adults and children is a large area in its own right and can be given only a brief mention here (for an introduction see Newcombe and Huttenlocher 2000). As indicated by the scattered mentions of findings with people, much contemporary work in this area is closely integrated with that on other species, especially in looking at spatial behavior in terms of a number of distinct subprocesses and in failing to find evidence for overall cognitive maps. Nowhere is this more evident than in a prominent opinion piece titled "Human spatial representation: Insights from animals" (Wang and Spelke 2002). Wang and Spelke proposed that rather than depending on an enduring allocentric map, much human spatial behavior depends on momentary egocentric representations, specifically dead reckoning, orienting by the geometry of surrounding space, and viewpoint-dependent matching of remembered to current views of the environment. Evidence for each of these processes comes from animal data like that reviewed throughout this chapter and from analogous experiments with people. In one key example, people viewed a room with a few objects in it and were then blindfolded, disoriented, and asked to point to the objects and the corners of the room. Errors in pointing indicated that the objects had not been integrated either into a map of the room as a whole or into a single configuration (Wang and Spelke 2000). Evidence that recognition of a familiar scene takes longer from a novel viewpoint supports the suggestion that encoding is viewpoint-dependent. However, more recent research (Burgess 2006) indicates that human spatial representation has both egocentric and allocentric components, which exist in parallel. In experiments like those just summarized, greater experience, a larger environment, and other factors make allocentric representations more evident. This approach is clearly much in the spirit of other research emphasized in the present chapter in dissecting spatial cognition into distinct parallel but interacting mechanisms and eschewing discussion of overall maps. Whether two systems defined in terms of function, egocentric and allocentric, will provide a useful way forward remains to be seen.

8.4.5 Conclusions

The cognitive map has been seducing investigators for over 60 years, but perhaps it is no more than a metaphor based on human introspection. Just as with theory of mind and other hypothetical mechanisms discussed later in this book, attributing cognitive mapping to an animal may be an unwarranted exercise of anthropomorphism, and one that is not even very useful in explaining human behavior. Translating such an intuitively appealing explanation of apparently intelligent behavior into testable implications in a way that researchers agree on is never easy. When the results of behavioral tests cause theorists to revise ambiguous and slippery concepts, agreement can become almost impossible. In the case of cognitive mapping, there is little if any unambiguous evidence that any creature gets around using a representation that corresponds to an overall metric survey map of its environment. The exceptional cases in which animals satisfy one or another classic criterion for mapping-like behavior by taking novel short cuts in the absence of direct cues from the goal (Singer, Abroms, and Zentall 2006; W. Roberts, Cruz, and Tremblay 2007) or finding their way home when displaced (R. Menzel. et al. 2005) are better explained by reference to what cues the animals are actually using, how they are using them, and how they come to do so than to the ill-defined notion of a cognitive map.

8.5 Summary

The study of spatial orientation is a very active area using a wide variety of species and approaches from fieldwork to neuroscience (Box 8.4). Among areas of research in comparative cognition it is exemplary, perhaps unique, in the way in which data and theorizing have been integrated across species and approaches as for example in the book edited by Jeffery (2003). The richest bodies of data come from three very different groups of animals: small nocturnal rodents (rats and hamsters), diurnal, central-place foraging insects (bees, wasps, and ants), and birds that orient over tens to hundreds of kilometers (homing pigeons and migratory species). The ways in which these animals perceive the world (consider for instance the very different visual systems of rats, pigeons, and bees) and the cues relevant for orientation in their natural environments differ enormously, yet some orientation mechanisms such as landmark learning or path integration and their interactions have been analyzed in a way that cuts across phyla. To some extent, this integrative approach has resulted in a theoretical orientation based on ideas from human psychology being replaced by one rooted in data from nonhuman animals.

Box 8.4 Space in the Brain

The study of what parts of the brain, particularly in mammals, help to control spatial behavior and how they do so is a vast area of contemporary behavioral neuroscience. The fact that the hippocampus is important for spatial memory in both mammals and birds has already been alluded to in Chapters 2 and 7, but in itself this does not tell us much about how brains actually represent space. Until recently the primary relevant information consisted of evidence for *place cells* in the rat hippocampus, single cells that fire when the rat is in a particular location within a laboratory enclosure. However, although the properties of place cells have been studied in some depth, one property seemingly essential for coding space is apparently lacking, topographic organization. That is, cells close together in the hippocampus do not

ars, but perhaps it is
with theory of mind
ttributing cognitive
opomorphism, and
Translating such an
avior into testable
When the results of
oncepts, agreement
there is little if any
representation that
it. The exceptional
ı for mapping-like
ues from the goal
remblay 2007) or
re better explained
are using them, and
ve map.

riety of species and
reas of research in
in which data and
for example in the
e from three very
amsters), diurnal,
at orient over tens
ies). The ways in
ery different visual
ientation in their
echanisms such as
been analyzed in a
h has resulted in a
ng replaced by one

l spatial behavior
The fact that the
has already been
bout how brains
sisted of evidence
is in a particular
place cells have
ace is apparently
ocampus do not

necessarily fire to places close together in space. Moreover, the same cell may have a place field (i.e., area in which it is active) in more than one enclosure (cf. Jeffery 2003).

The last few years have seen major advances in understanding how space is actually coded in the brain (see McNaughton et al. 2006). A major discovery is cells in the entorhinal cortex, *grid cells,* that map space in a periodic pattern whose spatial scale increases in an orderly way across layers of the medial entorhinal cortex. Combined with signals from cells sensitive to the animal's head direction and perhaps self-motion cues, these have the potential to code changes in an animal's position. Just as a unique time could in principle be coded by simultaneously reading the states of multiple oscillators with different frequencies (Section 9.3.2), so a location in space can be represented in terms of overlapping tessellations of tiles ranging from quite small to nearly the size of the space. New paradigms that require rats to navigate by dead reckoning have been used to probe the function of grid cells. In addition, hippocampal place cell recordings from rats moving over much longer distances than in traditional studies (e.g., 18 meters) show that the size of place fields increases across the hippocampus, perhaps providing a means to encode both details of space and general spatial context (see Hasselmo 2008).

This chapter began with descriptions of the wide range of mechanisms animals use for getting around. By itself, each of them has advantages and disadvantages. Dead reckoning is most useful for short journeys back and forth to a central place, especially in an environment with relatively few landmarks, as in the dark or on the desert. Other ways of getting back and forth to a starting place include route learning both in the sense of a memorized sequence of motor patterns (response learning) and in the sense of a sequence of responses to landmarks. Dead reckoning and route learning in either sense leave the animal lost if it is displaced too far off its usual route. However, stimulus generalization between familiar and unfamiliar views of the environment gives route learning some flexibility.

The varieties of spatial information—from landmarks, beacons, dead reckoning, environmental shape—are processed in different cognitive modules which take different kinds of input and output decisions about what distance and/or direction to move relative to different kinds of cues. This raises the question of how the outputs of different spatial modules are combined during the acquisition and use of spatial information. Are different kinds of information processed in parallel, do they compete for learning as in conditioning, or are they integrated in some other way? When are modules used in a hierarchical manner, and why? When spatial cues have acquired their significance, do they compete for control or are their outputs averaged? When does each kind of combination rule operate? For instance, does the system that has been more reliable during evolution or individual experience or that evolved earlier take precedence? A great deal of attention has been devoted to the question of whether any animal integrates different sources of spatial information into a unified allocentric representation of distances and directions, a cognitive map. This question turns out to be difficult to answer, partly because cognitive maps can mean different things to different people. Focusing on the specific cues available to animals and how they are used in specific situations provides better understanding of how animals get around than attempting to prove or disprove use of a cognitive map.

Further reading

Useful recent reviews of most aspects of spatial cognition can be found in the books edited by Jeffery (2003) and Wasserman and Zentall (2006b) and the online "cyberbook" edited by Brown and Cook (2006). Boinski and Garber's (2000) *On the Move* emphasizes field studies of animal movement patterns; the July 2007, special issue of *Animal Cognition* discusses numerous examples from primates. Newcombe and Huttenlocher's (2000) book is an excellent introduction to the development of spatial cognition in children very much in the same spirit as this chapter. The book by Wallraff (2005), the review by Bingman and Cheng (2005), and the special section of the August 11, 2006, issue of *Science* provide more information on homing and migration.

Some of the classics in the area are still well worth reading. These include the first six chapters of Gallistel (1990), with all aspects of animal spatial cognition discussed in the context of human navigation. The first two chapters of O'Keefe and Nadel's (1978) book are an excellent introduction to philosophical and psychological notions about space. A facsimile of this entire book is available free at http://www.cognitivemap.net. For a discussion of exploration, Berlyne (1960) still contains a lot of wisdom and a summary of much psychological and ethological literature.