

A connectionist model of path integration with and without a representation of distance to the starting point

ROLAND MAURER

University of Geneva, Geneva, Switzerland

Path integration (dead reckoning) is the process by which animals use self-generated movement information to continuously update a representation of their position relative to some starting point. This allows them to return to this point, even in the absence of exteroceptive positional information. Path integration is a bicoordinate process: Both distance and direction to (or from) the starting point have to be maintained in the representation. Recent theoretical work based on neurobiological observations seems to lead to the conclusion that, in rodents at least, path integration is map based, implying a geocentered (earthbound) coordinate system. The present paper discusses, as a word of caution, the opposite notion that path integration in an impoverished, single-coordinate egocentered (animal-bound) polar coordinate system could account for some rather puzzling behavioral data previously recorded on hamsters. This admittedly strange position hinges on observations made on connectionist models of path integration that were developed to simulate the homing behavior of hamsters.

Should the reader enter a large, fully darkened hall, closing the door before venturing (perhaps following a somewhat convoluted path) for some distance into the hall without any visual information, he or she would nonetheless be able to return by a straight path and with fair precision toward the starting point of that trip in darkness—namely, the door. The process which lies behind this feat is called *path integration* (M.-L. Mittelstaedt & H. Mittelstaedt, 1980).

Path integration, the capacity to return to the origin of a path even in the absence of exteroceptive information about location, has long been observed (Darwin, 1873) and, more recently, has been studied experimentally in humans (see, e.g., Fujita, Klatzky, Loomis, & Golledge, 1993; Sauv e, 1989). But, in spite of the cognitive abilities that it seems to imply (formally it is equivalent to vector addition and hence requires trigonometric computations), this ability is not limited to human beings: It also lies at the core of short-distance navigation in animals. Species for which path integration has been quantitatively studied include (but are not limited to) spiders (see, e.g., G rner, 1958), desert ants (see, e.g., M ller & Wehner, 1988), bees (see, e.g., Bisetzky, 1957), hamsters (see, e.g., Etienne, Maurer, & Saucy, 1988; S guinot, Maurer, & Etienne, 1993), and dogs (S guinot, Cattet, &

Benhamou, in press). In the return-to-start (or similar) tasks studied, those widely different species shared a common trait: After L-shaped outward journeys, the return journey tended, in many cases, to deviate toward the first leg of the outward journey, suggesting that perhaps a globally similar mechanism was involved. This sometimes led to unwarranted generalizations of models proposed for one species to other species (see Maurer & S guinot, 1995).

Path integration can be described as a computational process that allows the animal to keep track of its position relative to the starting point of its trip by resorting to information about its own movements. The animal can then return to that starting point, even in the absence of exteroceptive spatial information such as landmarks. Path integration, however, is not restricted to a single function: Beyond giving the animal the means to localize the origin of its journey, it may also be used to keep a constant direction in space (as in the phenomenon of turn alternation or angle compensation; see, e.g., Barnwell, 1965; Merkel & Fischer-Klein, 1973), to identify specific points in ambiguous environments (see, e.g., Etienne, Maurer, & S guinot, 1996; Hoffmann, 1985; Margules & Gallistel, 1988), to keep a search path centered on a point of maximal probability (see, e.g., Bovet, 1984; Hoffmann, 1984; M ller & Wehner, 1994; Wehner & Srinivasan, 1981), to continuously point toward a target discontinuously visible (see, e.g., Hill, 1979; and, in humans, Fukusima, Loomis, & Da Silva, 1997; Loarer, 1989), or to continuously localize two points in space (Etienne, Hurni, Maurer, & S guinot, 1991).

Path integration is a very pervasive process that occurs continuously, even when visible landmarks are present, and can be used for orienting (see, e.g., Etienne, Joris-

This research was supported by Swiss NSF Grants 31-26227.89 and 31-39311.93. I thank A. Etienne for continued encouragement, Z. Schreter for introducing me into the realm of connectionism, V. S guinot for providing invaluable experimental data, and two anonymous reviewers for helpful comments on a previous version. Correspondence concerning this article should be addressed to R. Maurer, Laboratoire d'Ethologie, 54 route des Acacias, 1227 Carouge, Switzerland (e-mail: maurerr@uni2a.unige.ch).

Lambert, Maurer, Reverdin, & Sitbon, 1995). This indicates that this process is not merely a back-up system: It probably combines with all modes of spatial information and collaborates in structuring the long-term spatial knowledge of the animal (Etienne et al., 1996; Gallistel, 1990; Gallistel & Cramer, 1996; McNaughton et al., 1996; O'Keefe & Nadel, 1978). However, path integration in itself is not a memory process but rather a computational process: It only stores in short-term memory the very volatile intermediate results of a computation that has to be repeated at every step.

Path integration is similar in nature to dead reckoning, the process by which sailors once kept track of their position when navigating on uncharted or open seas (see Kayton, 1990). Because of that similarity, path integration in an animal context is also sometimes referred to as dead reckoning (see, e.g., McNaughton, Chen, & Markus, 1991). The idea that animals could resort to a mechanism akin to dead reckoning dates back at least to Darwin (1873); continuous integration of path elements was suggested by Murphy (1873).

In its simplest realization, path integration requires that only one spatial relation be continuously updated during locomotion—that which links the position of the animal to the position of the starting point of the trip. If the coordinate system is geocentered (earthbound), the coordinates of the animal will change as a result of locomotion; on the contrary, if the coordinate system is egocentered (animal-bound), the coordinates of the starting point will change.

In order to update this representation, the animal must measure the translational and directional components of each new "step" it makes and use these values and some algorithm to change the old representation into a new, updated one.

Because of its algorithmic, iterative nature, path integration in animals is prone to the accumulation of random errors. In addition, in many cases path integration seems to involve, not only random errors, but also systematic errors. Although the latter have long been overlooked (see Maurer & Séguinot, 1995), they provide an open window into the workings of the process: They are a tool that allows exploration of the nature of the algorithm embodied in path integration, in that they provide the means to test models of the algorithm.

Models of Path Integration

Despite its obvious importance in short-range navigation and in related tasks (see above)—which also implies that it is probably involved in the construction of long-term spatial representations, including cognitive maps (Gallistel, 1990; McNaughton et al., 1996)—path integration was a neglected subject for a number of years. Recently, however, it has been revived in a number of modeling attempts. Although most of them are strictly mathematical in nature, the latest such efforts aim at bridging the gap between the geometry of the problem and the neurobiological structure that addresses it. A brief account of these attempts is presented below (for a more detailed review of some of those models, see Mau-

rer & Séguinot, 1995; see also Benhamou & Séguinot, 1995, for a mathematical treatment).

M.-L. Mittelstaedt and H. Mittelstaedt (1980), who coined the term *path integration*, used the exact trigonometric description of the computation required as a model of animal path integration: The animal would perform a trigonometric decomposition of the elements of its path, and the resulting components would be summed up in order to yield the positional vector. In this model, the variables are defined in a geocentered coordinate system. The model was strongly embedded in the Von Neumann computer metaphor of the sixties and is mostly a cybernetic (flowchart) description of the exact trigonometric solution to the problem of integrating the elements of one's path. According to this model, systematic errors should occur only in the presence of conflicting directional information or (in the later version of M.-L. Mittelstaedt and Glasauer, 1991) of unusual locomotion speeds.

Benhamou, Sauv e, and Bovet (1990) also assumed trigonometric computations but within an animal-centered reference system. In this probabilistic model, the animal could only estimate its movement, and the estimations were affected by errors drawn from Gaussian distributions. The authors used this model to demonstrate that random noise in the assessment of step length causes a smaller scatter in the estimated locations of the starting point than does random noise in the assessment of rotation and also that the presence of an external, stable directional reference (i.e., a compass) to help measure the rotations decreases the scatter in the homing vectors.

Fujita, Loomis, Klatzky, and Golledge (1990) also resorted to trigonometry in an animal-centered coordinate system. However, they simplified the equations and eliminated from the computations variables whose effect was small. They also replaced sine and cosine functions by linear approximations. Later (Fujita et al., 1993), they assumed that the errors they observed in human subjects were due to linear deforming functions applied to the measurement of rotations and length of path segments.

S guinot et al. (in press) explained the directional bias observed in dogs after L-shaped outward journeys along lines somewhat similar to those of Fujita et al. (1993). They too resorted to trigonometry within an animal-centered coordinate system, with the additional assumption that the inferred motion parallax of the starting point (i.e., the angular change in the head-referred return direction) was overestimated by a constant factor.

M ller and Wehner (1988) altogether discarded trigonometry as the core process and developed an arithmetic model, similar to Jander's (1957) seminal attempt, in which path integration is essentially a weighted average of angles: The successive changes in direction, relative to the vector computed so far, are weighted by the reciprocal of the distance to the starting point and averaged together. Although the variables of this model can be expressed in egocentered coordinates (Benhamou & S guinot, 1995), in the authors' own formulation—despite Wehner's (e.g., Hartmann & Wehner, 1995; Wehner, Michel, & Antonsen, 1996) repeated use of the word *egocentric* in later de-

scriptions of the model—these variables are inscribed in a geocentered polar coordinate system. The model, devised to describe L-shaped and U-shaped journeys of *Cataglyphis* ants, yields a prediction about the homing error that is strikingly close to the observed error.

In the more recent theoretical framework of *neural networks*, Maurer and Schreter (1990) and Maurer (1993) investigated how *simplifying* neural models (see Sejnowski, Koch, & Churchland, 1988) could be put to use in the context of path integration. The present paper reports and extends some of the results of Maurer (1993).

In the same framework, Hartmann and Wehner (1995) have devised a model of path integration in ants that relies on a highly constrained architecture, involving dedicated structures for representing coordinates and for computing goniometric functions (e.g., angular differences). This model implements in neuronal form the model of Müller and Wehner (1988) and generates the same errors (hence, it also is a good rendition of the behavior of ants). Perhaps its most interesting contribution is the observation that minor parametric modifications allow the very same architecture to compute exact (trigonometric) path integration, refuting once and for all the idea that trigonometry is so computationally difficult for neuronal machinery that it must be replaced by an approximate solution.

Also in the neural network framework, but based on increasingly detailed neuroanatomical and neurophysiological data obtained on rats, McNaughton and colleagues have gradually developed over the years a very ambitious modeling effort (for a review, see McNaughton et al., 1996), which, in its most recent embodiment (Samsonovich & McNaughton, 1997), aims at demonstrating that the hippocampus and its place cells (O'Keefe & Nadel, 1978) form, together with related structures, a prewired synaptic matrix where long-term spatial relations are stored and which, at the same time, provides the computational substrate for geocentered path integration. This *map-based path integrator* is at variance with all other models of path integration (*naive path integrators*, in the words of Samsonovich & McNaughton, 1997), including the one described in the present paper, which postulate a dedicated, environment-independent structure performing the computations of path integration, only interacting with the *map* part of the system when needed (e.g., when a position has to be written on the map or when the integrator needs correcting because of its drift).

Other models (see, e.g., Touretzky & Redish, 1996; Wan, Touretzky, & Redish, 1994) also try to include path integration as a building block of spatial knowledge but do not provide an explicit model of how the computation—the integration itself—is carried out.

Representation of the Variables of Path Integration

Some of the models mentioned postulate that path integration is realized within a geocentered coordinate system. Gallistel (1990), analyzing path integration in his all-

encompassing review of representation, also favored earthbound Cartesian coordinates, because of computational weaknesses associated with polar coordinate computations (feedback loops cause the errors, unavoidable in biological systems, to grow exponentially).

However, alternative polar representations (see, e.g., H. Mittelstaedt, 1983) may overcome those drawbacks. Moreover, earthbound Cartesian coordinates seem an unnecessary assumption for a mechanism strongly egocentric both in its inputs (body movements, directions to a beacon or compass direction) and in its outputs (control of movements required to orient toward the target). Finally, given the general occurrence of path integration, this process must be a very primitive function of organisms, which arose early in phylogeny. Hence, it could well be based on early spatial-processing nervous and sensory structures, which exploit spatial information in an organism-centered reference system (see Schöne, 1984, pp. 34 ff). One of the aims of the present paper is to test this assumption in a neural network context.

As mentioned before, the models presented here are not based on detailed neurophysiological data (as opposed to those developed by McNaughton and colleagues), and, accordingly, their purpose is not to gather new knowledge about the workings of the neuronal structures embodying path integration. Their scope is strictly limited to the following questions: Can egocentered polar coordinate path integration (deemed computationally intractable in biological organisms by Gallistel, 1990) be realized by means of neural-like primitives? Can the outcome of some experiments (in particular, those of Séguinot et al., 1993) that seem to imply that hamsters make only a partial computation of their path be explained by the properties of such a computational device?

Definitions

One term and two abbreviations will be used repeatedly throughout this text and may require definition.

The *target* is that point in space whose position is updated in the animal's egocentered representation by means of path integration. In the present context, this corresponds to the place to which the animal will have to return after its outward journey—that is, the origin or starting point of that journey.

Cw and *ccw* stand for clockwise and counterclockwise, respectively.

GENERAL ARCHITECTURE

The models belong to the class of multilayer feedforward networks trained by backpropagation (see Rumelhart, Hinton, & Williams, 1986). The reason why simple multilayer feedforward networks were selected for the present models lies in the fact that there is as yet no firm knowledge regarding the neural substrate of path integration in mammals. Although the neural architecture accounting for path integration in mammals is slowly relinquishing its secrets, even the most advanced research

on the topic of the underlying architecture (McNaughton et al., 1996; Samsonovich & McNaughton, 1997) remains speculative.

I therefore chose to resort to a general class of networks that are based on a training rule that, although biologically implausible if taken at face value (Crick, 1989; Sharkey, 1990), may nonetheless converge toward solutions similar to those realized in central nervous systems (Crick, 1989; Lockery, Wittenberg, Kristan, & Cottrell, 1989; Zipser & Andersen, 1988).

Since multilayer feedforward networks are universal approximators (Hornik, Stinchcombe, & White, 1989), any such (sufficiently large) network should be able to compute the function(s) underlying path integration. However, the nature of the task and the imbedded constraints (in particular, those stemming from the architecture chosen and from the size of the network, which may be suboptimal) will lead the network being trained toward a particular solution. A study of such a solution may then provide information pertaining to the nature of path integration as seen from the vantage point of this particular neural computing device.

In order to apply the general architecture of multilayer feedforward networks to the problem at hand, I followed a method similar to Sharp's (1991), whereby, on the one hand, I used neurobiologically plausible assumptions to define how the activity patterns in the network would represent the positional and movement information, and, on the other hand, I built a simulated environment and a simulated animal in it whose behavior would indicate whether the simulation attempt had been successful. The characteristics of the chosen environment, the simulated animal, and the task it was to perform were directly inspired by a set of experiments carried out on hamsters in our lab.

MODEL I

Path Integration Within an Animal-Centered Polar Coordinate System

At the core of the model lies a single layer of units, some of them representing the position of the target relative to the animal and others encoding information about the movement of the animal. The represented position of the target must be updated when the animal moves. This is realized by having both classes of input (regarding position and movement) converge into a computational structure (two hidden layers, which provide computational power to the network) and produce a common output in an output layer, which then represents the new position of the target, based on its old position and on the movement (see Figure 1). Specific details about the representation chosen can be found in the following sections.

Units Encoding Movement Information

Three units encode the characteristics of the current step. For simplicity, it shall be assumed that the information on the current step is obtained through the efference

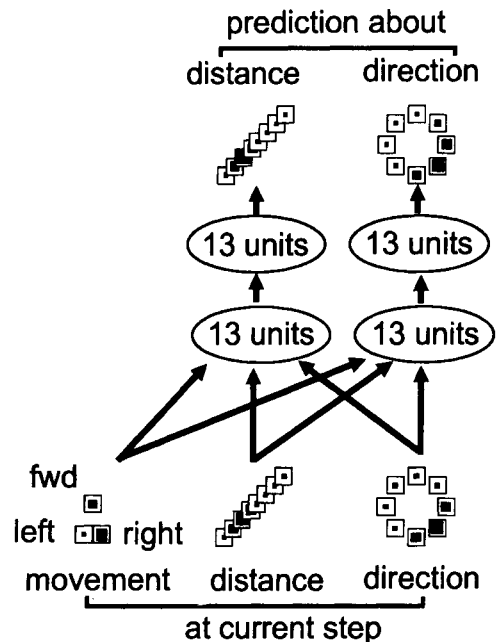


Figure 1. Model I, embodying path integration in polar, animal-centered coordinates. In the *lower* (input) layer, three units (at left) encode the values for the current step (step length, angle of left or right rotation), eight units (center) encode the current distance to the target (starting point of the outward journey) in a distributed fashion (units are displayed on a diagonal scale: the higher the point of maximal activation, the larger the distance), and eight units (at right) encode the current direction to the target, also in a distributed fashion (units are displayed on a compass rose, with *forward* toward the top of the figure). The shape of the groups of units is for legibility only and has no meaning or effect in the network; neither has the order of the units. The *upper* (output) layer contains two groups of eight units, whose pattern of activity should anticipate, respectively, the distance and direction to the target after the current step has been executed. The two *inner* layers give computational power to the network. The arrows show the pattern of connections between groups of units: in groups linked by arrows, every unit in the origin group is connected to every unit in the destination group. Above the first inner layer, the pathways for distance and direction are segregated. To perform path integration on more than one step, the network is *folded over itself*—that is, output values for step n are used as input values in step $n + 1$. At the time shown, the target lies behind and to the right of the simulated animal, which is about to perform a step with both forward movement and a rotation to the right (i.e., cw).

copy (von Holst & H. Mittelstaedt, 1950) of the motor commands issued from the motor centers rather than through proprioceptive or exteroceptive information. One unit stands for forward movement, the other two for (respectively) left (ccw) and right (cw) rotation. This dissociation of movement into translational and rotational components reflects observations of a similar dissociation into polar components in the nervous circuitry controlling head movement in the barn owl *Tyto alba* (Masino & Knudsen, 1990) and in the cat (Glassman, 1983), as well as whole-body movement in the frog *Rana pipiens* (Kostyk & Grobstein, 1987; Masino & Grobstein, 1989)

and in the rat (Chen, Lin, Green, Barnes, & McNaughton, 1994a, 1994b). All three units have linear responses proportional to the magnitude of the rotation and translation during the current step (for details, see the Appendix).

Units Encoding One Polar Coordinate

Each of the two polar, animal-centered coordinates of the target (direction and distance) is encoded in a group of eight units. (Any number beyond two may have been used; the choice of eight results from a compromise between precision of encoding, as tested in pilot experiments, and computational speed.) Every unit encodes one *preferred*, or nominal, value of the coordinate, to which it maximally responds (e.g., one of the directional units maximally responds when the origin is at 0 or *straight forward*, another when the origin is 45° to the left, and so forth; one of the distance units maximally responds when the target is at distance 0, another when it lies at distance 100, and so forth). The response function of every unit is triangular and decreases linearly as the value of the coordinate moves away from the preferred value, until the response reaches a resting (minimal) level (for details see the appendix). Because of the broad sensitivity range of the units, there will always be more than one of them (normally, three to four for each coordinate) activated above resting level.

The preferred values of the units are chosen so as to provide regular coverage of the whole range of directions and distances required (e.g., the preferred directions of the direction units are distributed evenly over 360°).

This *value encoding* (Ballard, 1986, 1987), spread over a number of units with large, overlapping sensitivity ranges, is a very efficient way of representing a continuous variable with arbitrary precision in a limited number of units. This encoding scheme is well-known as *coarse coding* (Hinton, McClelland, & Rumelhart, 1986). Coarse population coding, moreover, is relatively insensitive to noise in the response of single neurons (Sejnowski, 1988). Examples of such encoding in the nervous systems of animals are very common. For example, in the auditory map of the bat, the distance to the prey is encoded in distributed fashion: it generates an activity peak with a radius of 1,500 neurons (Baldi & Heiligenberg, 1988).

The triangular response curve of the units reflects that observed for the head direction cells in the rat's post-subiculum (Taube, Muller, & Ranck, 1990) and anterior thalamic nuclei (Taube, 1995): Those cells respond maximally when the rat's head lies in a particular direction in the horizontal plane, and their response decreases almost linearly to the resting level over approximately 30°–70° when the head direction changes (i.e., their sensitivity range covers some 60° to 140°). Finally, the isotropic distribution of preferred directions reflects the isotropic distribution of the preferred directions observed by the same authors for the head direction cells. The regularly spaced preferred distances of distance units and their triangular response curve were chosen by analogy.

Thus, units for direction and units for distance were similar in all respects, except, of course, that the activation pattern in the former would change as the direction to the target changed, whereas the activation pattern in the latter would change as a function of the distance between the animal and its starting point.

Units of the Computational Multilayer Structure

The output responses of the movement units and polar coordinate units converge on a two-layer computational structure (hidden layers) that provides computational power to the network. The response of each unit in this structure is a nonlinear, sigmoid function of the sum of its inputs (see the appendix), similar in its shape to the gamma distribution of the integral of the interval between neuronal spikes (see Hanson & Burr, 1990, p. 475). Its two 26-unit layers are actually divided into two half-layers each. Beyond the first layer, where they are merged together, the signals pertaining to distance are segregated from those pertaining to direction. This architecture was chosen in order to reduce the crosstalk between the two subtasks (see Jacobs, Jordan, & Barto, 1991; McClelland, 1986) by partially isolating one coordinate from the other (see Plaut, Nowlan, & Hinton, 1986). In pilot experiments (Maurer, 1993), this architecture yielded better results than did globally similar architectures where the signals were not kept separate.

Output Units

The output layer of the network, as mentioned, represents (predicts) the new position of the animal after the step has been completed. Hence, this layer replicates the input layer, except that it lacks the three units encoding the movement of the animal. The units of this layer, like the units of the hidden layers, respond only as a function of the activation arriving through their connections, and their transfer function is also sigmoid.

Training of the Networks

Training of the network is realized in a step-by-step procedure by means of backpropagation (Rumelhart et al., 1986). The simulated "hamster" is first randomly displaced to a point in a disk of 500-cm radius around the origin. This radius was chosen so that the system may be trained for the full range of distances (0–500 cm) a hamster could experience between any two points in the artificial environments of our lab (the largest dimension is the diagonal of the square arena mentioned below, equal to 500 cm). Next, the body axis of the simulated animal is randomly oriented; the units representing position in the input layer are set to the correct values of activation for that position, and the units representing movement in the same layer are set to the values corresponding to the current, randomly chosen step. Now, activation is forward-propagated through the network, the current step is executed, and the output of the network is compared to the correct values of activation corresponding to the new

position. The errors are then propagated back into the network, and the connections are adjusted. As hinted before, this “learning” procedure is not meant to simulate any specific learning process in the animal but to see how the problem at hand is decomposed and solved by a neural-like computing device. To cite Kohonen (1988, p. 250), “Learning in this context means improvement of the system performance, in the sense of some criterion, relating to use.” However, backpropagation is a special case of the Robbins-Monro stochastic approximation procedure (Robbins & Monro, 1951; White, 1989), and it will extract statistical information from the problem space, which is what an adapted organism is supposedly able to do.

Ten replicates of the network, with connections set to different initial random weights, were run 400 times each over a data set of 2,500 different position–movement pairs. The learning rate and the momentum term were chosen as 0.1 and 0.9, respectively, following pilot experiments. After this initial training, the residual error of each of the trained replicates was assessed by running them for another 250 random steps and computing the average error in direction and distance over those 250 steps. The best replicate was then trained a further 5,600 passes over the same data set. This lengthy procedure was chosen because it created a smooth learning curve toward an asymptote that was the performance limit of the network, while at the same time minimizing the chances of getting caught in an early local minimum.

During one simulation cycle, the activation pattern in the position units (encoding the target location in animal-centered coordinates) and the activation pattern in the movement units (encoding the change in orientation and the linear displacement of the current step) are propagated forward in the network. The units in the output layer will be activated as a result of this, and their pattern will correspond to the network’s estimation of the position of the target (direction and distance to the animal) after the current step. This is formally equivalent to a one-step-ahead prediction (Weigend, Huberman, & Rumelhart, 1990). What is needed, however, for continuous path integration is a dynamic memory loop yielding an iterated *prediction*. To that end, the trained network is *folded*, and the position units of the output layer and of the input layer become one. That is, the output of the network is fed back to the network before the next step, then combined with the movement input for that step. The network, after an initialization to the correct positional values, can then compute iteratively the position of the target on the basis of the movement information alone. This embodies the integration process into a dynamic memory loop, similar to Jordan’s (1986) sequential (recurrent) network architecture or to Kamath and Keller’s (1976) model of the integration in the vestibulo-ocular reflex. Similar ideas for a dynamic memory loop can be found, for example, in Reeke, Sporns, and Edelman (1990) and can also be traced back to the pioneering work of McCulloch and Pitts (1943). Little is known about such memory loops in

the animal. Mormyrid fish, however, seem to rely on such a dynamic memory loop to update electrosensory information (Heiligenberg, 1991). It has also been suggested (Blair & Sharp, 1995) that a thalamocortical circuit may integrate angular head movement to compute geocentered head direction in much the same kind of memory loop.

Comparing the Performance of the Networks With That of Hamsters

The performance of the trained network, which translates into the capacity to correctly update its representation of the position of the starting point, was evaluated by replicating, in a simulated environment, two series of experiments on real hamsters (see Etienne et al., 1988; Séguinot et al., 1993). In these experiments, the hamster (*Mesocricetus auratus* W.) inhabited a round (diameter, 220 cm) or square (side, 355 cm) arena and established its quarters and its granary in a constantly accessible nest peripheral (round arena) or internal (square arena) to the arena. During each experimental trial, the hamster was induced to leave the nest and follow a bait until it reached a pile of food (hazelnuts). The hamster then pouched the food and started its return path toward the nest-granary, where it deposited the food. The experiments took place in a fully darkened room, the bait being made visible to the hamster by a very dim light-emitting diode located on it. To the fully adapted human eye, the light emitted by this diode did not even allow the hamster to be seen, much less the surrounding environment. Infrared lighting in the room and an infrared camera located centrally above the arena allowed the animal’s movements to be recorded. Other precautions were taken to prevent the animal from using acoustical, olfactory, or tactile cues. The animal’s use of geomagnetic cues had been ruled out by a previous series of experiments (for all experimental details and procedures, please refer to the original papers).

In one series of experiments (round arena; Etienne et al., 1988), the hamster was made to walk a circular path around the food (three full revolutions) before being given access to it. In the second series of experiments (square arena; Séguinot et al., 1993), the hamster had to follow an L-shaped outward journey from the nest to the food. These experiments were considered for simulation because the former lays stress on the computation of rotations and the latter on the computation of distances.

I built a simulated animal able to move in a two-dimensional space that was featureless except for a starting point (the “nest”), which could be “perceived” only while the simulated animal was in contact with it. The animal could be positioned anywhere in the simulated environment, it could be made to follow specific paths, and it could also be directly piloted by the output of the network.

The physical dimensions are obviously arbitrary in a simulated environment. Therefore, to allow easy comparisons, the relevant elements (environment size, animal size, nest size, path size, translational and rotational

step sizes) were scaled in proportion to their laboratory counterparts.

To perform the simulations, videotaped hamster paths were digitized and transformed into sequences of appropriate "motor" commands for the simulated animal. Thus, in the simulated environment, the paths followed by the animal were identical to the paths imposed on the real animals by having them follow baits. The digitization steps (and hence the steps made by the simulated animal) correspond to a maximum length of 3 cm in the real environment, about equal to the stride length of slowly moving hamsters, as shown by tracks left in the sawdust substrate of the arenas.

The correct direction and distance to the target were initialized in the network when the animal "left its nest." Thus, the network direction was initialized to "straight behind," and the network distance was initialized to the distance between the center of the nest and the point where the animal lost contact with it (10 cm in the round arena and 20 cm in the square arena). After initialization, the only external inputs received by the network were the activations of the movement units, corresponding to the digitized path and possibly different for every step. At each step, the activations in the output layer were copied back to the corresponding part of the input layer (dynamic memory loop). At the end of the outward journey, a vector was computed from the pattern of activations in the output layer (see the appendix). This vector originated at the final position of the animal and ended at the target position computed by the network. The differences in direction and length between this vector and the correct position vector were computed. Since the real animals' estimate of direction could be observed but not their estimate of distance (unless hamsters are cued by some signal—e.g., olfactory markings on the ground—they usually do not stop in their tracks after covering the correct distance to the goal to which they have been trained; see Etienne et al., 1998; Séguinot et al., 1993), only the directional errors of the network are reported here and compared by means of Hotelling's confidence ellipse (Batschelet, 1981) to the corresponding values found for the real hamsters.

Results

The first notable result of the training process was that the folded network, as expected, came to embody a dynamic memory loop for each coordinate. In both groups of units, a single, stable *bump* of unit activity (distributed over two to four neighboring units) appeared and remained present during the iteration of the network. In effect, the network, once folded, generated two activity bumps or *packets* (see Samsonovich & McNaughton, 1997), each representing the current value of one of the coordinates. As was the case in Samsonovich and McNaughton's simulations, the activity bumps resulted directly from the dynamic properties of the whole system, since they appeared over the course of a few iterations even if the network was initialized with random activity




Figure 2. The shape of an outward journey involving three clockwise revolutions plus an additional 90° arc, as digitized from an actual videotaped hamster track (for clarity, the different revolutions have been slightly separated). There were four types of journeys with either cw or ccw *main* rotations and either cw or ccw additional 90° arcs. This was meant to prevent the hamster's learning to associate the final position of its body axis with the direction to the nest. For the analysis, however, the journeys with cw and ccw additional 90° arcs were collapsed together. Correspondingly, there were four types of simulated outward journeys, with differing final parts, and the vectors for cw and for ccw final parts were averaged, as was the case with real hamsters.

(in this case the coordinate values represented by the bumps were, of course, arbitrary). In the absence of any movement input, the bumps slowly drifted away from their initial positions.

Three Revolutions

Figure 2 shows the shape of one videotaped outward journey (one straight leg, three revolutions clockwise around the arena center, followed by an additional 90° arc, a left turn of 90°, and finally a straight leg to the center of the arena). Figure 3 shows the average return directions of 6 hamsters, each of them tested 15 times in cw and 15 times in ccw rotations (the letters identify the different individuals). Figure 4 shows the directions computed by the trained network for the same outward journeys. It can be seen that, despite the small number of processing elements, the network computes the return direction with a precision similar to that of hamsters after a three-revolution outward journey. In one direction of turning, the rotations of the simulated animal are very precisely summed up, and the vector correctly points within a few degrees of home. In the other direction, the summation is not as good, but the resulting error remains close to the largest error made by the real animals.

This asymmetry in the computations cannot be easily explained. Although this is very improbable, the asymmetry may be due to some asymmetry during training or to some bias in the initial random distribution of weights of the network. On the other hand, the asymmetry may indicate that the network has difficulties in optimizing all the facets of the computation at the same time and that




Figure 3. The average return directions of 6 real hamsters, each of them tested 15 times in clockwise rotations (hatched arrows), and 15 times in counterclockwise rotations (solid arrows). The letters identify the different individuals. Data are from Etienne et al. (1988).




Figure 4. Simulation runs. The direction to the target as computed by network of Model I at the end of the four different circular outward journeys. The simulated animal was made to follow the different tracks (hatched lines). Upon losing contact with the nest (rectangle to the left of figure), the network was initialized to the correct values of activation corresponding to the animal's position. It then computed the homing vector iteratively, on the basis of movement information only. At the end of each outward journey, the homing vector was computed from the pattern of activations in each group of coordinate units and printed (thin radial lines). CW (resp. CCW) indicates the return vectors computed at the end of the two journeys involving cw (resp. ccw) revolutions. Compare with the real-animal data of Figure 3.

it used its limited power to solve only part of the task (namely, updating for ccw turns rather than for cw turns). Similar, unexplained asymmetries can be observed in the responses of individual hamsters (see, for instance, animals *b* and *d* in Figure 3).

For the statistical assessment of differences between the orientation of the simulated and real animals, it was necessary to get rid of the asymmetries. This was done by averaging values corresponding to cw and ccw simulated outward journeys (after the required sign changes). The resulting mean was compared to the average homing vectors of the sample of hamsters by means of Hotelling's confidence ellipse. For journeys involving a rotation around the arena center, the difference was found to be nonsignificant (see Table 1, column Model I).

L Shapes

Figure 5 shows the (idealized) shape of all L-shaped outward journeys and the corresponding average return directions of 5 hamsters, each tested three times for each type of outward journey (except for the shortest L shapes, which were replicated in another part of the arena, giving rise to 2 x 3 measures).

Figure 6 shows the corresponding directions computed by the network. It can be seen that, in this case also, rather strong asymmetries affect the computed homing vector as a function of the shape of the outward journey. In the left part of the figure, while the animal proceeds along the second leg of the journeys, the network seems to update the direction of the homing vector in a rather regular fashion, although not fast enough to keep it aligned with the target. In the right part of the figure, corresponding to a clockwise forced turn, only two vectors out of three seem to have been rotated during the second leg of the journey. Again, this could be taken as an indication that the task was difficult for the network, which possibly oscillated between the different variables it tried to incorporate into the update function.

In a two-coordinate path integrator, errors in one coordinate will necessarily cause subsequent errors in the other coordinate. There was a small correlation ($r = .453$) between absolute errors in distance and absolute errors in direction in the 10 situations in which Model I was tested. This suggested that direction might be computed better without any representation of distance in the network. The idea is that, in some cases, it may be better to have no information at all than to have erroneous infor-

Table 1
Orientation Error (in Degrees) of Return Vectors of the Network of Model I (Bicoordinate, Trained Over a 500-cm Disc), Model II (No Distance Coordinate, Trained Over a 500-cm Disc), and Model IIb (No Distance Coordinate, Trained Over a 200-cm Disc), Compared With the Orientation of Real Hamsters

Journey	Hamsters	Model I	Model II	Model IIb
3 rev.	25.0	25.7‡	8.1*	32.4‡
1 m-1 m	11.5	23.4‡	-19.0‡	0.7‡
1 m-2 m	0.2	-13.5*	-18.3‡	3.5‡
2 m-1 m	20.2	-12.9‡	-3.2*	16.1‡

Note—The data for cw and ccw forced turns have been collapsed together after the required sign changes. The significance of the difference has been assessed by means of Hotelling's confidence ellipse. * $p < .05$. † $p < .01$. ‡*n.s.*, $p > .05$.




Figure 5. The six different L-shaped outward journeys (idealized shapes), and the corresponding average return directions of real hamsters. The nest is shown to scale as a large circle. The animal was led along the L (i.e., first from the bottom of the figure toward the top, then along a horizontal line) by means of a baited spoon to the food located at the end of the second leg (black dots). The length of the segments is 1 m. The arrows show the average return directions of a group of 5 hamsters (for the two 1-m-1-m situations, each animal was tested twice, in different parts of the arena—hence the double arrow). Each animal ran three times every type of journey. Data are from Séguinot et al. (1993).

mation. To test this quite counterintuitive hypothesis, another model network was devised, which contained no units representing the distance to the target.

MODEL II

Path Integration Without a Representation of Distance

In contrast to Model I, Model II (see Figure 7) was completely devoid of the two eight-unit groups (one in the input layer, the other in the output layer) representing the distance to the target. As the required task was simpler (only one polar coordinate to update; no crosstalk), fewer computational units were used in the hidden layers, in order to speed up the simulations. The number of units in the hidden layers (15 per layer) was chosen so that the network would total half the number of connections of Model I. In all other respects, this model was similar to Model I. In particular, the training took place in exactly the same way.

Results

Three Revolutions

Figure 8 shows the average directions computed by the network after an outward journey comprising three cw or ccw revolutions. It can be seen that after a three-revolution outward journey, the network computes the return direction with a precision significantly better (see Table 1, column Model II) than that of the hamsters, despite its very small number of processing elements. In the present case, the absence of distance units is of little consequence, since the outward journey is mostly circular and ends at a point aligned with the original direction of the animal exiting from the nest; so, the geometric problem of returning home could be solved by keeping a count of rotations only (equiv-

alent to maintaining an inertial—or other—compass). It is, then, hardly surprising that the simulated animal keeps its homing vector within a few degrees of the home direction.

L Shapes

The case of L-shaped outward journeys is more interesting. In this case, the second leg of the journey is orthogonal to the direction of the animal when it leaves the nest, and this geometry would fool an attempt to return by means of compasslike information only. Contrary to what may be expected considering that the network was built without a representation for distance, it can be plainly seen in Figure 9 that the network does a fair job of progressively correcting the direction of the target as the simulated animal proceeds along the second leg of the journey.

The overall precision of the computation performed by the networks in these different types of outward journey is then roughly comparable to that performed by real hamsters, despite the very small number of neural-like processing units involved and an architecture missing any form of representation of distance (homing errors in the range of 3°–19°, as compared with that of 0°–25° for hamsters). Each single situation, however, showed significant differences, indicating that, besides the superficial similarity, the network followed different rules (see Table 1, column Model II). Still, it must be emphasized that the direction to origin *is* updated as the animal moves orthogonally to the first leg of the journey, although there is now no explicit representation of the distance to the target in the network.

A closer analysis of the simulation runs reveals that, whenever the simulated animal is following the second




Figure 6. Simulation runs. The direction to the target as computed by network of Model I at the end of each type of L-shaped outward journey. The simulated animal was made to follow six different L-shaped tracks (hatched lines). Upon losing contact with the nest (large circle toward bottom of figure), the network was initialized to the correct values of activation corresponding to the animal's position. It then computed the homing vector iteratively on the basis of movement information only. At the end of each outward journey (small circles), the home vector was computed from the pattern of activations in each group of coordinate units and printed. Compare with the real-animal data of Figure 5.

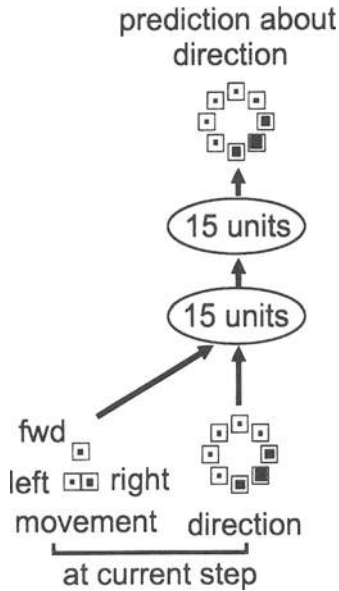


Figure 7. Model II, a network without units representing distance to target. Since it contains only direction units, the network has only one pathway, and its total number of connections is set to half that of Model I. For further explanations, see Figure 1 and text.

leg of a journey, it behaves as if it were trying to update the direction to the target as a function of its progression and of a single, fixed distance, somehow built into the network. A simple geometrical estimation allows us to reckon that distance: the intersection of the return vectors with the *y*-axis (collinear with the first leg of the journey) lies at a distance of, on average, 250 cm from the second segments of the corresponding outward paths (those segments that are horizontal on the figures). This implies that the network has recorded, in the process of training, a single distance that is something like the average of all distances it had been exposed to during the training phase. As was said before, during training, the simulated animal was placed randomly and with a *constant density* over a disk of 500-cm radius around the target. The average distance to the target will then be 333 cm and *not* 250 cm. Therefore, the information the network seems to have extracted from the situation is the median of the range of distances to which it was submitted.

In order to verify this, the best replicate of the same network was trained again, with the same procedure, except that the training set corresponded to positions within a 200-cm radius (2.5 times shorter than before). In this case, the median of the range of distances was, of course, 100 cm. After training, the network (which will be referred to as Model IIb) piloted the animal, using the same rule as before (update vector along the second leg, without taking heed of the actual length of the first leg), but, this time, the fixed distance learned by the network was obviously 2.5 times shorter than before (see Figure 10). It may be noted that, with a radius set to 200 cm, the surface (12.6 m³) covered by the training set was equal to

the surface of the square arena used in hamster experiments involving L-shaped journeys.

If one compares the results of the simulation, as shown in Figure 10, with the real-hamster data of Figure 5, two things are obvious. First and foremost, the striking similarity between the network's performance and the hamster's (no significant differences; see Table 1, column Model IIb) suggests that the hamsters' orientation observed in this particular situation may well be explained by a process similar to the one at work in this network.

By and large, the simulated animals compute the return vector by taking into account the distance they are made to walk along the second leg of the journey, whereas the distance along the first leg is not actually taken into account but is replaced by a fixed distance of about 1 m. This translates into the fact that the return directions after 1-m-1-m journeys are correct, but that after 2-m-1-m journeys, they are roughly parallel to the former and hence incorrect. If one assumes that a similar process underlies the orientation of the hamsters, what is the fixed distance learned by the animals? If one applies to the real hamster data the same geometrical analysis that was carried out on the simulated returns, one finds that the animals seem to have recorded a fixed distance of slightly less than 1 m, very close to the data of Figure 10. Could it be that, at least in the experimental situation described here, the animal somehow evaluates the surface of its artificial home range and computes a *best-fit* distance on that basis? Or does it perhaps record the range of possible distances during its training? Hamsters have to be familiarized with the experimental task before the actual experiments can begin, so they may have the opportunity to extract the required information at that moment.

Although such a device would be an elegant and economical way to solve the problem of orienting by path in-




Figure 8. Simulation runs. The direction to the target as computed by the network of Model II (without a representation of distance) at the end of the four different circular outward journeys (thin radial lines). For further explanations, see Figure 4. Compare to the real-animal data of Figure 3.




Figure 9. Simulation runs: the direction to the target as computed by the network of Model II (without a representation of distance) at the end of each type of L-shaped outward journey. For further explanations, see Figure 6. Compare with the real-animal data of Figure 5.




Figure 10. Simulated L-shaped outward journeys for the network of Model IIb (without a representation of distance), trained over a smaller disk, equivalent in surface to the square arena. The figure shows the direction to the target, as computed by the network when it has reached the end of every type of L-shaped outward journey. For further explanations, see Figure 6. Compare with the real-animal data of Figure 5.

tegration in our reduced experimental setting (where imprecise returns cause no negative consequences, except that the animal will then take slightly longer to reach the nest), the critical experiments to prove this post hoc hypothesis have not yet been done.

DISCUSSION

The models presented here resort to a simplified but fairly realistic subset of general neural properties (activation of units as a function of inputs, propagation of activation through connections, divergence and convergence of connections, a rule-extracting training process). These neural properties are coupled to more specific ideas about possible underlying architectures (distributed, nontopographic representation of spatial variables and recurrent connections) in order to realize path integration in an egocentered polar coordinate representation. The simulations show that such an architecture can generate a dynamic memory for spatial position. Iteration of the networks on the basis of movement information alone—that is, without correction from the environment—produced a running computation of a home vector, roughly similar to path

integration as observed in hamsters tested in total darkness, whereby the values encoded in the dynamic memory changed in accordance with the movement of the animal.

The dynamic memory of the trained system took, as expected, the form of self-sustained activity bumps in each of the polar coordinate unit sets. As is the case in Samsonovich and McNaughton (1997), the activity bumps stemmed directly from the global dynamics of the system, contrary to Hartmann and Wehner (1995), where the activity bumps (actually not bumps but sectors) result from local connections between units. Whereas Hartmann and Wehner had to resort to a highly constrained architecture in order to implement the various elements of path integration, the models presented here have minimal architectural constraints as regards the connections between units. In particular, Hartmann and Wehner assumed that vector information is encoded in neural maps where anatomically neighboring neurons represent similar values and the anatomical proximity of the neurons influences the way they are connected. Such vector maps may exist in ants, but they have yet to be found. In contrast, connection schemes independent of neuronal proximity are known to exist in rodents, where both the head direction cells and the place cells represent space in nontopographic fashion (see McNaughton et al., 1996).

With those indications in favor of a geocentered representation for space in rodents (especially as evidenced by the existence of head direction cells, which provide the nervous system with the equivalent of an inertial—hence geocentered—compass), one may wonder why an egocentered representation was chosen for the present models. Indeed, in light of the current knowledge about the neuroanatomy and neurophysiology of rodents, it would seem that the most plausible candidate for a path integration scheme would be the *map-based path integration* (MPI) class of models, as advocated by Samsonovich and McNaughton (1997), since all the necessary building blocks have been identified in the hippocampus and related structures. Such an MPI scheme implies a geocentered coordinate system. However, the differences between an MPI and a *naive path integration* scheme may be subtle and, therefore, hard to distinguish by behavioral tests.

Moreover, although this is unlikely, it might be that the head direction cells actually encode the egocentered direction to the starting point of path integration, but in a context where the animal thinks it is very far away from this starting point. One could, for instance, assume that the animal sets its path integrator to zero in its home cage when it is taken out of it. Then, when the rat is passively carried to the experimental setup located in another room, it may be induced to update its path-integrated position on the basis of a large distance value, in effect transforming egocentered direction to the nest into a geocentered compass within the experimental space. The difference is subtle and may have been overlooked. An easy way to test this would be to carry the animal to two identical re-

cording rooms devoid of orienting cues, located in opposite directions with respect to the home cage. If the head direction cells code for geocentered head direction, their response will be roughly parallel in the two environments; if, on the contrary, they encode egocentered home direction, they should respond with a 180° phase shift.

At any rate, path integration as behaviorally observed in real animals does not easily fit into the current MPI picture. Neither spiders, nor ants, nor bees, nor hamsters, nor dogs, nor humans behave as though their path integrator were affected by simple drift. It is only in theoretical models (see, e.g., Benhamou et al., 1990) that the path integrator behaves as if it were affected by a random drift; in practice, it does not.

Besides these theoretical considerations, perhaps the most striking (and counterintuitive) result of those simulations is that the model closely replicated the behavior of hamsters (progressively updating the homing vector along the second leg of L-shaped journeys) only when it had been deprived of the possibility of explicitly representing the distance to the target. In addition, changing the training procedure so that the area over which the network was trained reflected the overall area (not linear dimension) of the real hamsters' arena again improved the fit of the model to the observed data.

Although, as mentioned above, a number of species in unrelated taxa show patterns of errors after L-shaped outward journeys that are similar to those reported here for hamsters (they also behave as if they underestimate the first leg of the journey), care should be taken not to confuse the issue. If all of these results were to be explained in terms of the hypothesis presented here, such errors would mean that, in every experiment, the constant value used by the path integrator was shorter than the actual length of the first segment of the path. Although there are not many such experiments, it is unlikely (although not impossible) that this was the case in all of them. More to the point, ants (one of the species tested) are known to measure distance fairly precisely, mainly by means of optic flow (Ronacher & Wehner, 1995), and to use it in path integration (a striking example can be found in Wehner, 1982, Figure 7). A model of path integration without distance representation would hardly make sense for this species. Moreover, the errors observed in ants have satisfactorily been explained (Müller & Wehner, 1988) by assuming that distance and direction are both represented but are updated by an approximate algorithm. However, as pointed out by Maurer and Séguinot (1995), the errors observed in *mammals* do not follow exactly the same scheme and need a different explanation. The present simulations only purport to account for path integration as observed in hamsters (and possibly for other mammals).

There are some indications, beyond the data considered here, that hamsters and other mammals may sometimes resort to simplified rules for the evaluation of distance in path integration. In a series of experiments that also took place in darkness (Etienne et al., 1988), the ham-

ster was lifted in a box above the arena floor immediately upon exiting from its nest. The arena was then shifted under the box, until the arena center was under the box. The box was then turned by 90°, and lowered onto the arena surface, and the animal was released. In these conditions, the outward journey was limited to the walk between the nest and the end wall of the box, so path integration should have indicated a very short distance (some 20 cm). The animal could then have been expected to slow down after it had covered that distance. In reality, however, all of the animals tested ran the full distance to the arena wall without the slightest hesitation. Thus, they obviously had learned that the usual distance they were required to run was 1 m or more. It may be concluded that they kept a running record of the rotations needed to compensate for the rotations experienced in the box but did not keep, or at least did not use, a record of the outward distance they moved.

Perhaps more to the point, golden hamsters (as mentioned above) need to be cued by external stimuli in order to stop at a goal toward which they are homing by means of path integration. This was shown by experiments (Etienne et al., 1998) in which the hamster was led by different indirect routes (made of a peripheral section followed by a partial crossing of the round arena) to a goal (some food buried in the sawdust substrate) located halfway between the periphery and the arena center. The training proceeded until the animal could go unassisted from the periphery to the goal. This training phase was expected to provide the animal with a memorized nest-to-goal vector, which the animal could then reel up through path integration in order to zero in toward the goal. In transfer trials, the food was removed, and the arena was cleansed, in order to erase possible olfactory traces where the food had been. The animal was then led along a new peripheral route and released. In many cases, the hamster chose the correct direction and, hence, actually walked over the place where the food should have been. This indicates that path integration could indeed provide the animal with rather accurate directional information. In almost no case, however, did the animals actually stop at (or search around) the correct location. As soon as some olfactory or tactile information was added (e.g., by the experimenter's touching the sawdust substrate), the animals stopped and searched. This need for an external confirmation could be a direct consequence of the absence of accurate distance information postulated in the present models.

There are examples on a much larger scale. Bovet (1984) showed that gray squirrels passively displaced way beyond their home range are able, in some conditions at least, to orient toward the home range. However, after they have covered a distance roughly equivalent to the size of their home range, they turn back and return to the release point, from which they initiate another search path in another direction. This shows that they know something about the *average* or *maximal* distance that makes sense for a given home range, and it may well be that this

average distance is the only one they use in their large-scale path integration mechanism.

The use of a fixed instead of a variable distance has the obvious drawback that it diminishes the precision of the computation of direction for all cases in which the orthogonal leg of the journey is not at that fixed distance from the target. However, it also has advantages: (1) it considerably reduces the complexity of the required computations, and (2) it completely eliminates the accumulation of errors in one variable, since that variable is eliminated.

There are simple ways of radically improving the performance of a system based on the principles discussed here. For instance, instead of using a single stored distance, it could resort to a set of distances (i.e., a set of functions for directional update), each applied when needed. For example, in a natural setting, the animal may estimate the distance to its nest on the basis of the time it has spent out of the nest or on some measure of the novelty of the surroundings and apply the corresponding update function.

If distance is only registered as a single value or as a limited set of values within the computational system, one may expect to observe effects of scale in experiments on path integration. A 5-m–5-m L-shaped outward journey, for a given subject, may, for instance, not be equivalent to a 50-m–50-m one. As argued in Maurer and Séguinot (1995), path integration in mammals is still a poorly understood phenomenon, and data about it remain scarce, especially concerning the effect of the scale over which path integration is performed.

Finally, the hypothesis developed here—namely, that computational accuracy in path integration may be obtained at the expense of accuracy in the representation of distance (since this variable is replaced by a static approximation)—also brings to light the fact that path integration, if it only relies on an average value of distance obtained by sizing up the environment in some way, will be affected by the animal's experience, instead of being an immutable process. To the best of my knowledge, only Samsonovich and McNaughton (1997) have also suggested this, but in another context (according to their map-based path integrator scheme, path integration should be less accurate in a novel environment, because of the progressive recruiting of map cells). As a corollary, the computations carried out by the path integrator may be affected not only by the shape of the outward journey but also by the geometry of the experimental setup and the history of the animal's experience with it. The large errors observed in some experiments may result from a complex interaction between these variables and may essentially reflect the unnaturalness of the experimental situation, rather than the inherent shortcomings of path integration.

REFERENCES

- BALDI, P., & HEILIGENBERG, W. (1988). How sensory maps could enhance resolution through ordered arrangements of broadly tuned receivers. *Biological Cybernetics*, **59**, 313-318.
- BALLARD, D. H. (1986). Cortical connections and parallel processing: Structure and function. *Behavioral & Brain Sciences*, **9**, 67-120.
- BALLARD, D. H. (1987). Interpolation coding: A representation for numbers in neural networks. *Biological Cybernetics*, **57**, 387-402.
- BARNWELL, F. H. (1965). An angle sense in the orientation of a millipede. *Biological Bulletin*, **128**, 33-50.
- BATSCHLEET, E. (1981). *Circular statistics in biology*. London: Academic Press.
- BENHAMOU, S., SAUVÉ, J.-P., & BOVET, P. (1990). Spatial memory in large-scale movements: Efficiency and limitation of the egocentric coding process. *Journal of Theoretical Biology*, **145**, 1-12.
- BENHAMOU, S., & SÉGUINOT, V. (1995). How to find one's way in the labyrinth of path integration models. *Journal of Theoretical Biology*, **174**, 463-466.
- BISSETZKY, A. R. (1957). Die Tänze der Bienen nach einem Fussweg zum Futterplatz [The dance of bees after they walk to the feeding station]. *Zeitschrift für vergleichende Physiologie*, **40**, 264-288.
- BLAIR, H. T., & SHARP, P. E. (1995). Anticipatory head direction signals in anterior thalamus: Evidence for a thalamocortical circuit that integrates angular head motion to compute head direction. *Journal of Neuroscience*, **15**, 6260-6270.
- BOVET, J. (1984). Strategies of homing behavior in the red squirrel, *Tamiasciurus hudsonicus*. *Behavioral Ecology & Sociobiology*, **16**, 81-88.
- CHEN, L. L., LIN, L. H., GREEN, E. J., BARNES, C. A., & McNAUGHTON, B. L. (1994a). Head-direction cells in the rat posterior cortex: I. Anatomical distribution and behavioral modulation. *Experimental Brain Research*, **101**, 8-23.
- CHEN, L. L., LIN, L. H., GREEN, E. J., BARNES, C. A., & McNAUGHTON, B. L. (1994b). Head-direction cells in the rat posterior cortex: II. Contributions of visual and idiothetic information to the directional firing. *Experimental Brain Research*, **101**, 24-34.
- CRICK, F. (1989). The recent excitement about neural networks. *Nature*, **337**, 129-132.
- DARWIN, C. (1873). Origin of certain instincts. *Nature*, **7**, 417-418.
- ETIENNE, A. S., HURNI, C., MAURER, R., & SÉGUINOT, V. (1991). Two-fold path integration during hoarding in the golden hamster? *Ethology, Ecology, & Evolution*, **3**, 1-11.
- ETIENNE, A. S., JORIS-LAMBERT, S., MAURER, R., REVERDIN, B., & SITBON, S. (1995). Optimizing distal landmarks: Horizontal versus vertical structures and relation to background. *Behavioural Brain Research*, **68**, 103-116.
- ETIENNE, A. S., MAURER, R., BERLIE, J., DERIVAZ, V., GEORGAKOPOULOS, J., GRIFFIN, A., & ROWE, T. (1998). Cooperation between dead reckoning (path integration) and external position cues. *Journal of Navigation*, **51**, 23-34.
- ETIENNE, A. S., MAURER, R., & SAUCY, F. (1988). Limitations in the assessment of path dependent information. *Behaviour*, **106**, 81-111.
- ETIENNE, A. S., MAURER, R., & SÉGUINOT, V. (1996). Path integration in mammals and its interaction with visual landmarks. *Journal of Experimental Biology*, **199**, 201-209.
- FUJITA, N., KLATZKY, R. L., LOOMIS, J. M., & GOLLEDGE, R. G. (1993). The encoding-error model of pathway completion without vision. *Geographical Analysis*, **25**, 295-314.
- FUJITA, N., LOOMIS, J. M., KLATZKY, R. L., & GOLLEDGE, R. G. (1990). A minimal representation for dead-reckoning navigation: Updating the homing vector. *Geographical Analysis*, **22**, 326-335.
- FUKUSIMA, S. S., LOOMIS, J. M., & DA SILVA, J. A. (1997). Visual perception of egocentric distance as assessed by triangulation. *Journal of Experimental Psychology: Human Perception & Performance*, **23**, 86-100.
- GALLISTEL, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- GALLISTEL, C. R., & CRAMER, A. E. (1996). Computations on metric maps in mammals: Getting oriented and choosing a multi-destination route. *Journal of Experimental Biology*, **199**, 211-217.
- GLASSMAN, R. B. (1983). Dissociation of vertical and horizontal components of somesthetic orientation-localization during recovery from cortical damage: Implications regarding central associative functions. *Physiological Psychology*, **11**, 47-53.
- GÖRNER, P. (1958). Die optische und kinästhetische Orientierung der Trichterspinnne *Agelena labyrinthica* [The optical and kinaesthetic orientation of the funnel-web spider *Agelena labyrinthica*]. *Zeitschrift für vergleichende Physiologie*, **41**, 111-153.

- HANSON, S. J., & BURR, D. J. (1990). What connectionist models learn: Learning and representation in connectionist networks. *Behavioral & Brain Sciences*, **13**, 471-518.
- HARTMANN, G., & WEHNER, R. (1995). The ant's path integration system: A neural architecture. *Biological Cybernetics*, **73**, 483-497.
- HEILIGENBERG, W. (1991). The neural basis of behavior: A neuroethological view. *Annual Review of Neuroscience*, **14**, 247-267.
- HILL, D. E. (1979). Orientation by jumping spiders of the genus *Phidippus* (Araneae: Salticidae) during the pursuit of prey. *Behavioral Ecology & Sociobiology*, **5**, 301-322.
- HINTON, G. E., MCCLELLAND, J. L., & RUMELHART, D. E. (1986). Distributed representations. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1: Foundations* (pp. 77-109). Cambridge, MA: MIT Press.
- HOFFMANN, G. (1984). Homing by systematic search. In D. Varjú D & H. U. Schnitzler (Eds.), *Localization and orientation in biology and engineering* (pp. 192-199). Berlin: Springer-Verlag.
- HOFFMANN, G. (1985). The influence of landmarks on the systematic search behaviour of the desert isopod *Hemilepistus reaumuri*: II. Problems with similar landmarks and their solution. *Behavioral Ecology & Sociobiology*, **17**, 335-348.
- HORNIK, K., STINCHCOMBE, M., & WHITE, H. (1989). Multilayer feed-forward networks are universal approximators. *Neural Networks*, **2**, 359-366.
- JACOBS, R. A., JORDAN, M. I., & BARTO, A. G. (1991). Task decomposition through competition in a modular connectionist architecture: The what and where vision tasks. *Cognitive Science*, **15**, 219-250.
- JANDER, R. (1957). Die optische Richtungsorientierung der Roten Waldameise (*Formica rufa* L.) [The optical directional orientation of the red ant (*Formica rufa* L.)]. *Zeitschrift für vergleichende Physiologie*, **40**, 162-238.
- JORDAN, M. I. (1986). *Serial order: A parallel distributed processing approach* (ICS Report 8604). San Diego: University of California, Institute for Cognitive Science.
- KAMATH, B. Y., & KELLER, E. L. (1976). A neurological integrator for the oculomotor control system. *Mathematical Biosciences*, **30**, 341-352.
- KAYTON, M. (1990). *Navigation: Land, sea, air and space*. New York: IEEE Press.
- KOHONEN, T. (1988). *Self-organization and associative memory* (2nd ed.). Berlin: Springer-Verlag.
- KOSTYK, S. K., & GROBSTEIN, P. (1987). Neuronal organization underlying visually elicited prey orienting in the frog: I. Effects of various unilateral lesions. *Neuroscience*, **21**, 41-55.
- LOARER, E. (1989). Un exemple de coopération vision-représentation: Le contrôle visuel intermittent des déplacements locomoteurs [An example of cooperation between vision and representation: Intermittent visual control of locomotory motion]. *Psychologie Française*, **34**, 55-62.
- LOCKERY, S. R., WITTENBERG, G., KRISTAN, W. B., JR., & COTTRELL, G. W. (1989). Function of identified interneurons in the leech elucidated using neural networks trained by back-propagation. *Nature*, **340**, 468-471.
- MARGULES, J., & GALLISTEL, C. R. (1988). Heading in the rat: Determination by environmental shape. *Animal Learning & Behavior*, **16**, 404-410.
- MASINO, T., & GROBSTEIN, P. (1989). The organization of descending tectofugal pathways underlying orienting in the frog, *Rana pipiens*: I. Lateralization, parcellation, and an intermediate spatial representation. *Experimental Brain Research*, **75**, 227-244.
- MASINO, T., & KNUDSEN, E. I. (1990). Horizontal and vertical components of head movement are controlled by distinct neural circuits in the barn owl. *Nature*, **345**, 434-437.
- MAURER, R. (1993). *L'intégration du chemin ou navigation à l'estime chez l'animal: Formalisation de modèles neuromimétiques* [Path integration or dead reckoning in living beings: Neural network models]. Unpublished doctoral dissertation, University of Geneva, Switzerland.
- MAURER, R., & SCHRETER, Z. (1990). Sensori-motor spatial learning in connectionist artificial organisms. In F. Heylighen, E. Rosseel, & F. Demeyere (Eds.), *Self-steering and cognition in complex systems: Toward a new cybernetics* (pp. 160-182). New York: Gordon & Breach.
- MAURER, R., & SÉGUINOT, V. (1995). What is modelling for? A critical review of the models of path integration. *Journal of Theoretical Biology*, **175**, 457-475.
- MCCLELLAND, J. L. (1986). The programmable blackboard model of reading. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2. Psychological and biological models* (pp. 122-169). Cambridge, MA: MIT Press.
- MCCULLOCH, W. S., & PITTS, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, **5**, 115-133.
- MCNAUGHTON, B. L., BARNES, C. A., GERRARD, J. L., GOTHARD, K., JUNG, M. W., KNIERIM, J. J., KUDRIMOTI, H., QIN, Y., SKAGGS, W. E., SUSTER, M., & WEAVER, K. L. (1996). Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *Journal of Experimental Biology*, **199**, 173-185.
- MCNAUGHTON, B. L., CHEN, L. L., & MARKUS, E. J. (1991). "Dead reckoning," landmark learning, and the sense of direction: A neurophysiological and computational hypothesis. *Journal of Cognitive Neuroscience*, **3**, 190-202.
- MERKEL, F. W., & FISCHER-KLEIN, K. (1973). Winkelkompensation bei Zwergwachteln (*Excalfactoria chinensis*) [Angle compensation in Indian blue quails (*Excalfactoria chinensis*)]. *Die Vogelwarte*, **27**, 39-50.
- MITTELSTAEDT, H. (1983). The role of multimodal convergence in homing by path integration. *Fortschritte der Zoologie*, **28**, 197-212.
- MITTELSTAEDT, M.-L., & GLASAUER, S. (1991). Idiopathic navigation in gerbils and humans. *Zoologische Jahrbücher: Abteilung für Allgemeine Zoologie und Physiologie der Tiere*, **95**, 427-435.
- MITTELSTAEDT, M.-L., & MITTELSTAEDT, H. (1980). Homing by path integration in a mammal. *Naturwissenschaften*, **67**, 566.
- MÜLLER, M., & WEHNER, R. (1988). Path integration in desert ants, *Cataglyphis fortis*. *Proceedings of the National Academy of Sciences*, **85**, 5287-5290.
- MÜLLER, M., & WEHNER, R. (1994). The hidden spiral: Systematic search and path integration in desert ants, *Cataglyphis fortis*. *Journal of Comparative Physiology A*, **175**, 525-530.
- MURPHY, J. J. (1873). A mechanical analogy. *Nature*, **7**, 483.
- O'KEEFE, J., & NADEL, L. (1978). *The hippocampus as a cognitive map*. Oxford: Oxford University Press, Clarendon Press.
- PLAUT, D. C., NOWLAN, S. J., & HINTON, G. E. (1986). *Experiments on learning by back propagation* (Tech. Rep. CMU-CS-86-126). Pittsburgh: Carnegie Mellon University, Computer Science Department.
- REEKE, G. H., SPORNS, O., & EDELMAN, G. M. (1990). Synthetic neural modeling: The "Darwin" series of recognition automata. *Proceedings of the IEEE*, **78**, 1498-1530.
- ROBBINS, H., & MONRO, S. (1951). A stochastic approximation method. *Annals of Mathematical Statistics*, **22**, 400-407.
- RONACHER, B., & WEHNER, R. (1995). Desert ants *Cataglyphis fortis* use self-induced optic flow to measure distances travelled. *Journal of Comparative Physiology A*, **177**, 21-27.
- RUMELHART, D. E., HINTON, G. E., & WILLIAMS, R. J. (1986). Learning internal representation by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1. Foundations* (pp. 318-362). Cambridge, MA: MIT Press.
- SAMSONOVICH, A., & MCNAUGHTON, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network. *Journal of Neuroscience*, **17**, 5900-5920.
- SAUVÉ, J.-P. (1989). *L'orientation spatiale: Formalisation d'un modèle de mémorisation égocentrée et expérimentation chez l'homme* [Spatial orientation: A model of egocentered memory, and experimentation on man]. Unpublished doctoral dissertation, Université d'Aix-Marseille II.
- SCHÖNE, H. (1984). *Spatial orientation*. Princeton: Princeton University Press.
- SÉGUINOT, V., CATTET, J., & BENHAMOU, S. (in press). Path integration in dogs. *Animal Behaviour*.
- SÉGUINOT, V., MAURER, R., & ETIENNE, A. S. (1993). Dead reckoning

in a small mammal: The evaluation of distance. *Journal of Comparative Physiology A*, **173**, 103-113.

SEJNOWSKI, T. J. (1988). Neural populations revealed. *Nature*, **332**, 308.

SEJNOWSKI, T. J., KOCH, C., & CHURCHLAND, P. S. (1988). Computational neuroscience. *Science*, **241**, 1299-1306.

SHARKEY, N. (1990). There is more to learning than meets the eye (or ear). *Behavioral & Brain Sciences*, **13**, 506-507.

SHARP, P. E. (1991). Computer simulation of hippocampal place cells. *Psychobiology*, **19**, 103-115.

TAUBE, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience*, **15**, 70-86.

TAUBE, J. S., MULLER, R. U., & RANCK, J. B. (1990). Head direction cells recorded from the postsubiculum in freely moving rats: I. Description and quantitative analysis. *Journal of Neuroscience*, **10**, 420-435.

TOURETZKY, D. S., & REDISH, A. D. (1996). Theory of rodent navigation based on interacting representations of space. *Hippocampus*, **6**, 247-270.

VON HOLST, E., & MITTELSTAEDT, H. (1950). Das Refferenzprinzip (Wechselwirkungen zwischen Zentralnervensystem und Peripherie) [The principal of refference (interaction between central nervous system and periphery)]. *Naturwissenschaften*, **37**, 464-476.

WAN, H. S., TOURETZKY, D. S., & REDISH, A. D. (1994). Towards a computational theory of rat navigation. In M. Mozer, P. Smolensky, D. Touretzky, J. Elman, & A. Weigend (Eds.), *Proceedings of the 1993 Connectionist Models Summer School* (pp. 11-19). Hillsdale, NJ: Erlbaum.

WEHNER, R. (1982). Himmelsnavigation bei Insekten: Neurophysiologie und Verhalten [Celestial navigation in insects: Neurophysiology and behavior]. *Neujahrsblatt der Naturforschenden Gesellschaft Zürich*, **184**, 1-132.

WEHNER, R., MICHEL, B., & ANTONSEN, P. (1996). Visual navigation in insects: Coupling of egocentric and geocentric information. *Journal of Experimental Biology*, **199**, 129-140.

WEHNER, R., & SRINIVASAN, M. V. (1981). Searching behaviour of desert ants, genus *Cataglyphis* (Formicidae, Hymenoptera). *Journal of Comparative Physiology*, **142**, 315-338.

WEIGEND, A. S., HUBERMAN, B. A., & RUMELHART, D. E. (1990). *Predicting the future: A connectionist approach* (Tech. rep.). Stanford: Stanford University, PDP Research Group.

WHITE, H. (1989, December). Neural-network learning and statistics. *AI Expert*, pp. 49-52.

ZIPSER, D., & ANDERSEN, R. A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, **331**, 679-684.

APPENDIX

Activation and Transfer Functions of Units

Activation of Units in Input Layer

Movement units. The activation of the movement units is defined as follows.

For the *Forward* unit: $a_F = 0.8 (s/s_{MAX}) + 0.1$, where s is equal to the current step length and s_{MAX} to the upper bound of step lengths—that is, 3 cm. Hence, the activation is limited to the interval [0.1 . . . 0.9]. This limitation was chosen to avoid boundary effects that were due to the characteristics of the sigmoid transfer function.

For the *Left Turn* and *Right Turn* units: $a_T = 0.8 \cdot d \cdot (r/r_{MAX}) + 0.1$, where r is equal to the current rotation (change in direction), r_{MAX} to the upper bound of rotations (12°), and $d = 1$ if the rotation takes place in the nominal direction of the unit (i.e., left for unit Left) and 0 otherwise. Here also, the activation is limited to the interval [0.1 . . . 0.9].

Units encoding polar one polar coordinate. *Direction* and *distance* units: for the i th unit in the group of units, its activation is defined as

$$a_i = \max (0, [0.8 \cdot (1 - (|v - V_i|/r_i))] + 0.1,$$

where v is the current value of direction (distance); V_i , the preferred value of the unit; and r_i , its sensitivity range. Within the range of sensitivity of the unit, the activation is a triangular function of the target's direction (distance) relative to the preferred value. It decreases linearly from a maximum of 0.9, when the target is in the preferred value of the unit, to a resting level of 0.1, when the target is 90° (200 cm) or more from the preferred direction.

Transfer Function of All Units Beyond the Input Layer

I used the classic logistical sigmoid transfer function,

$$o_j = 1/(1 + e^{-(net_j + theta_j)}),$$

where o_j is the output of the unit considered, net_j the sum of the inputs arriving to it, and $theta_j$ the input arriving from the bias unit (the latter having no other meaning than to provide in an elegant way a threshold value for each unit).

Decoding the Activation Pattern

In order to determine the value of the homing vector at any given step (for statistical or display purposes), it is necessary to compute the population vector for each of the two polar coordinate *populations* of units.

For direction units, this is done by computing the mean vector (i.e., by vectorially summing the preferred directions of the cells, each weighted by the current activation of the cell minus 0.1, then by dividing the sum by the number of units; the direction of the resulting mean vector is then taken to be the direction that the group is representing).

For distance units, the preferred values of the units are first proportionally mapped onto angular values (with a constant proportionality factor, 100 cm being mapped onto 10°, 200 cm onto 20°, and so on), the mean vector is computed as above, and then the resulting direction is transformed into a distance value by reverse mapping. The reader may satisfy himself that this procedure (chosen because it allows the same algorithm to compute the population vector for both types of units) indeed provides the distance represented within the population of units, whereas simply applying the weighted arithmetic mean does not.

(Manuscript received March 31, 1997;
revision accepted for publication October 20, 1997.)