Biomimetic Bayesian models of navigation: How are environment geometry-based and landmark-based strategies articulated in humans?

Julien Diard (Julien.Diard@upmf-grenoble.fr)
Laboratoire de Psychologie et NeuroCognition CNRS-UPMF
Grenoble, France

Panagiota Panagiotaki and Alain Berthoz
Laboratoire de la Physiologie de la Perception et de l'Action, Collège de France–CNRS,
Paris France

Abstract

We propose a computational model of human navigation, which encompasses both geometry-based and landmark-based navigation strategies. This model is based on a study of human cognitive strategies during a path memorization task in a Virtual Reality (VR) experiment. Participants were asked to memorize predefined paths in a large-scale virtual city (COSMOpoliS©). Our computational model qualitatively reproduces the results of this experiment. This model uses the Bayesian formalism, and focuses on the interplay between the elementary cognitive strategies hypothesized above. It offers an original interpretation of the way these strategies might be articulated, departing from the classical hierarchical structure. This novel view might be fruitful for robotic models from a biomimetic perspective, where managing representations of large-scale and complex environments is still a challenge.

Keywords: Bayesian modeling; human navigation; navigation strategies; landmark-based navigation; path integration.

We discuss here the results of an experiment, in which we have explored the existence of elementary cognitive strategies used for spatial encoding in humans. We have found that, while navigation mainly relies on landmark recognition and encoding for path memorization, the sudden disappearance of these brings in light a back-up mechanism strategy enabling the participants to navigate, although with less accuracy, using geometrical cues alone. These are the first evidence of equivalent components between humans and animals in this context.

In models of navigation, these observations of "back-up" mechanisms usually lead to modeling independent subsystems of navigation, and portraying them as hierarchically articulated. We believe this view of independent subsystems being hierarchically articulated to be simplistic, as it merely pushes back the problem of understanding how different sources of information are integrated in the central nervous system.

In this paper, we propose a probabilistic model that tackles this problem in an original manner. We develop a model of navigation, which, although composed of a single component, can mimic both landmark-based and geometry-based navigation strategies. Bayesian inference is the principle, which enables this single representation of the environment to give rise to several navigation strategies. The overall behavior of our model is dictated by the availability of sensory cues. When there are no uncertainties about the sensed landmarks, our model performs as landmark based navigation. On the other hand, when landmarks are not sensed, the model performs as geometry based navigation. In the following, the term "navigator" will refer to a simulated, imaginary participant that would navigate in our virtual city according to our mathematical model.

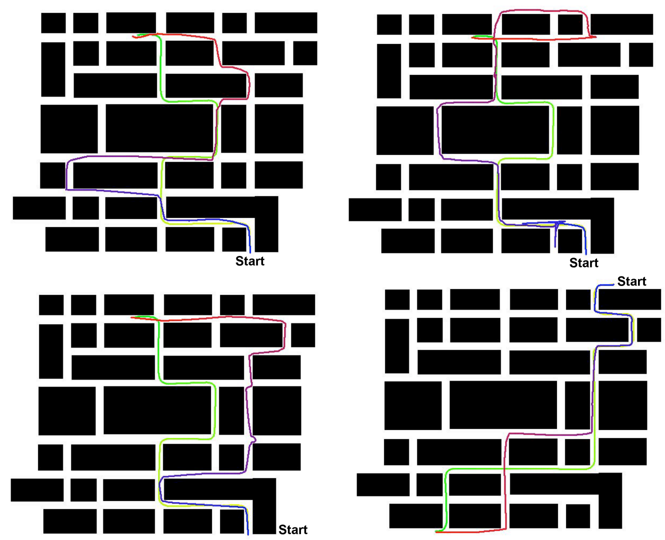

Figure 1: Top-view of the virtual city and archetypal errors observed in the condition where landmarks are removed between memorization and reproduction. In light gray (green), the learned paths. In black (blue to red), the reproduced paths by the participant. For example, in the top-right panel, note how the central building was passed from the right in the learned path, and from the left in the participant's reproduction.

The model qualitatively reproduces patterns of errors we observed in the COSMOpoliS© experiment. In this experiment, humans participants were immersed in a VR city using a VR helmet. They could navigate using a joystick for forward translations, and turn their body in the real world for virtual rotations (with a magnetic tracker set on the VR helmet). Participants were presented a movie of a trajectory, and were asked to memorize it (memorization phase). After seeing the movie twice, they were set in the starting end of the path and asked to reach its end, actively (reproduction phase). Landmarks (posters on walls, lampposts, etc.) were disposed in the city. Experimental conditions were defined by the availability of the landmarks in both the memorization and reproduction phase. We focus here in two conditions: in the Landmark condition, landmarks were available both in memorization and reproduction. In the Probe Trial condition, landmarks were removed between memorization and reproduction.

The data showed that all participants were able to successfully reach a goal in a virtual city in the Landmark condition. Data also showed that the goal was also reached when landmarks were removed between memorization and reproduction (Probe Trial). However, in that case, patterns of error could be observed quite frequently in the paths that participants made in order to reach the goal (see Fig. 1). Participants quite commonly reached the goal using a variant of the memorized path, passing buildings from the wrong side, for instance. Surprisingly, very few participants were actually conscious of these discrepancies.

The rest of this paper is structured as follows. Firstly, we briefly review the related work on hierarchical modeling of human, animal or robotic navigation. We then present the Bayesian model we developed in order to have our simulated navigator reproduce this pattern of error: we first introduce our simplifying assumptions, then define the model and describe its simulation. Finally, we discuss the interpretation of our model as was defined, as well as of the relevance of our simplifying assumptions. The paper concludes on a discussion on the way our assumptions could be relaxed, yielding perspectives on the future work.

Related modeling works

Both life sciences and robotics have made the modeling of navigation capabilities of autonomous entities a crucial point of research, and a wide variety of models already exists. We focus here on hierarchical models of navigation.

In the domain of mobile robotics, modeling the environment that a robot has to face, usually in the form of a map, is a crucial problem. The most promising approaches rely on the probability calculus, thanks to its capacity for handling incomplete models and uncertain information. These approaches include – but are far from limited to – Kalman Filters, Markov Localization models, (Partially and Fully) Observable Markov Decision Processes (POMDP and MDP), and Hidden Markov Models (see (Diard, Bessière, & Mazer, 2003) for a general introduction).

In this domain of probabilistic modeling for robotics, hierarchical solutions are currently flourishing. The more active domain in this regard is decision theoretic planning: one can find variants of MDPs that select automatically the partition of the statespace (see for instance (Hauskrecht, Meuleau, Kaelbling, Dean, & Boutilier, 1998)). Another class of approaches that rely on deterministic notions is based on the extraction of a graph from a probabilistic model, like for example a Markov Localization model (Thrun, 1998), or a MDP (Lane & Kaelbling, 2002).

However, the main philosophy used by these hierarchical approaches is to try to extract, from a very complex but intractable model, a hierarchy of smaller models Automatically selecting the right decomposition is of course a very difficult problem. Moreover, even obtaining in the first place the initial, complex model, is still a difficult challenge in the general case.

From a bio-mimetic perspective, it appears obvious that a global, complex, large-scale model is not the starting point of the acquisition of representations of space (B. J. Kuipers, 2000). Therefore, some robotic approaches, integrating insights from biology, rather start from low-level behaviors and representations, and then try to combine them so as to obtain large-scale representations (Diard & Bessière, 2008; B. J. Kuipers, 2000; B. Kuipers, Modayil, Beeson, MacMahon, & Savelli, 2004; Victorino & Rives, 2004). Indeed, the study of navigation capabilities in life sciences assumes right from the start of its analysis that navigation is hierarchical in nature, as can be easily assessed experimentally (Voicu, 2003).

The hierarchies of models proposed in some of these works (Trullier, Wiener, Berthoz, & Meyer, 1997; Franz & Mallot, 2000; B. J. Kuipers, 2000; B. Kuipers et al., 2004) have several aspects: they are hierarchies of increasing navigation skills, but also of increasing scale of the represented environment, of increasing time scale of the associated movements, and of increasing complexity of representations. This last aspect means that global topologic representations, which are simple, come at a lower level than global metric representations, which are arguably more complex to build and manipulate. This ordering stems from the general observation that animals that are able to use shortcuts and detours between two arbitrary encoded places (skills that require global metric models) are rather complex animals, like mammalians. These skills seem to be mostly absent from simpler animals, like invertebrates.

Works by Jacobs and Schenk go a step further, by proposing the Parallel Map Theory (PMT) (Jacobs & Schenk, 2003), in which a study of phylogenetically equivalent neuroanatomical areas across different species helps hypothesize common hierarchies of representations of space. In other words, they propose a model of how the different layers in the above theories might be implemented in the central nervous system.

Finally, Wang and Spelke (Wang & Spelke, 2002) assume three subsystems, two of which being a path integration (PI) and a view dependent place recognition system. These two, in the context of the current paper, can be seen as analogous of what we will denote as the environment geometry-based and landmark-based navigation systems, respectively.

However, to the best of our knowledge, the question of how different subsystems of a hierarchy of models can exchange information in a principled manner is still an open issue. In other words, most existing models of animal navigation describe hierarchies by identifying individual layers, but do not tackle the problem of how these layers are linked. They usually assume that a supervisor subsystem is responsible for selecting the interaction between individual components, but rarely describe the way this supervisor could work, or even discuss its plausibility (*e.g.* the reference frame selection subsystem of Redish & Touretzky (Redish & Touretzky, 1997)).

Our model precisely proposes an original articulation between a representation of a memorized path and resulting strategies of navigation.

Model

In this section, we develop a Bayesian model, which qualitatively reproduces the observed patterns of errors (see Fig. 1). Being preliminary, our model requires several simplifying assumptions that we describe first. We then describe how, given these assumptions, this single model is defined and used in order to simulate the navigator in the virtual city in both the Landmark and Probe Trial conditions. We finally discuss the similarity between the simulation and experimental data.

Simplifying assumptions

Our model requires two major assumptions: the first concerns the identification of orientations by the navigator; the second concerns the identification of landmarks.

Firstly, we assume that the navigator uses a global reference frame for orientations. This means that an estimate of the navigator's bearing with respect to some origin is available at every moment. Given this estimate, the navigator knows which direction it is currently going. This helps it classify elementary displacements according to the direction followed. This implies a separation between the estimation of orientations and the estimation of positions. Neuroanatomically, such a separation appears to be plausible: estimations of orientations might be grounded in head-direction cells (Stackman & Taube, 1997; Taube, 1998); estimations of positions might be grounded in place cells (Redish & Touretzky, 1997). However, to the best of our knowledge, such a separation is rarely present in robotic models, where, usually, the pose x, y, θ of the robot is considered, with similar mathematical treatment for position x, y and orientation θ.

In order to be used, this orientation reference frame does not need global sensory cues. Indeed, instead of being based on some external cue, the origin could be based on the starting orientation of the navigator (Berthoz et al., 1999).

We further assume that this global reference in orientation does not drift during the navigation of the path. Indeed, in COSMOpoliS© and in our simulation, all angles between streets are 90° angles, thus reducing risks of disorientation (drifting of the orientation reference frame). With these assumptions, in our model, we only need four possible orientations for the global reference frame. In the following, we denote "up" the starting direction, "down" the opposed direction, and "left" and "right" the two remaining directions.

Secondly, we assume that landmarks in the virtual city are all unique and easily recognizable. We assume they are placed at the intersections or decision points, as it has been shown that the relevance of a landmark to solving navigation tasks is explicitly encoded in the central nervous system (Janzen & Turennout, 2004). We further assume that landmarks can be used to recognize all intersections in the city without errors. These assumptions allow the model to include certainties (probabilities of 1) about the landmark and their recognition, when they are available.

Model definition

We now define a two-variable model.

The first variable, denoted L_t, is the location at time t, *i.e.* the intersection the navigator is in, as defined by the landmark appearing at this intersection. Assuming n different landmarks and intersections in the virtual city, $l_1, l_2, \ldots, l_n$, we thus define: $L_t = \{l_1, l_2, \ldots, l_n\}$. The second variable, denoted A, is the direction that should be followed at intersection l_t. According to our assumptions concerning the global orientation reference frame, we define A by $A = \{\mathrm{up}, \mathrm{left}, \mathrm{down}, \mathrm{right}\}$.

We thus define the joint distribution:

$$P(A, L_t) = P(L_t)P(A \mid L_t),$$

by applying Bayes rule. The first term, $P(L_t)$, is the likelihood to be in some intersection. We define this term by a uniform probability distribution: $P(L_t = l_t) = 1/n$. The second term, $P(A \mid L_t)$, represents probability distributions over directions to follow, given the identity of the intersection the navigator currently is at. We define this term by Conditional Probability Tables (CPT). We assume the navigator identifies these CPTs during the path memorization phase of the experiment. In other words, during path presentation, the navigator counts the number of times it went "up", "down", "left" and "right", and builds the CPTs that reflect these frequencies. The CPTs follow Laplace succession law distributions, which are similar to histograms, except that probabilities for unobserved cases are never zero [1].

There is one such CPT for each landmark seen along the memorized path. As we have assumed all landmarks to be unique, and assuming that the paths never pass twice in the same intersection (which is the case in the COSMOpoliS© experiment), these learned distributions are all of the same type: the probability is close to 1 for the actual direction followed along the path, and close to 0 for the three directions not followed.

[1] Laplace succession law probability distributions merge a prior distribution with observed data. The formula is $P(A = i \mid L_t = l_t) = \frac{n_i + w}{N + kw}$, with n_i the number of times a particular case i has been observed, N the total number of observations, k the size of the domain of the variable, and w a parameter which tunes the speed at which the initial uniform distribution is modified as the observations are collected. A Laplace succession law converges toward a histogram when the number of observations N is large. Assessing a biologically plausible weight w is an open question (out of the scope of this paper and experiment).

Model usage

Having learned the CPTs during the path presentations, the model is now fully defined. The joint distribution $P(A, L_t)$ is available to the navigator, and we describe here how it can be used to drive the navigator during the path reproduction, in the Landmark and Probe Trial conditions.

In both cases, the navigator must decide, at each intersection, the direction to follow in order to accurately reproduce the path it memorized. In the Landmark condition, intersections can be identified during the learning phase as well as during the reproduction phase. Therefore, when arriving at an intersection, the value l_t of the current intersection is available, and can be used to select the relevant probability distribution over actions $P(A \mid L_t = l_t)$. Once this distribution is selected, choosing the action with the highest probability value will lead the navigator along the memorized path, without errors.

Alternatively, the navigator can draw at random according to the memorized probability distribution. In this case, errors in the reproduction could occur, their frequency depending on the parameter w chosen for learning CPTs.

In the Probe Trial condition, landmarks are not available in the city anymore in the reproduction. Therefore, when arriving at an intersection, it is not possible for the navigator to know the value of l_t. However, using Bayesian inference, $P(A)$ can be computed: it is the probability distribution over actions to take at each intersection, without knowing the intersection identity. The computation is as follows:

$$P(A) = \sum_{L_t} P(A, L_t) \propto \sum_{L_t} P(A \mid L_t).$$

This computation yields the best estimate available to the navigator in order to choose what direction to go at each intersection during the reproduction. Drawing at random according to $P(A)$ allows the navigator to reproduce the memorized path in the best manner, given the absence of visual cues.

Model simulation

We have simulated the model in an idealized version of the VR city, abstracting ourselves from issues related to the small scale of the VR city COSMOpoliS©. We call this ideal, simulated city EQUApoliS. EQUApoliS is a regular, infinite grid of square blocks, with simulated unique landmarks at each intersection.

We defined a path to be learned and reproduced (narrow black path in Fig. 2). When the navigator memorizes this path, it learns by observation CPTs, one for each intersection. One such CPT is shown Fig. 3.

In the Landmark condition, the memorized path is accurately reproduced. Indeed, recall that the CPTs in this condition can be read so as to provide, at each intersection, a probability distribution over actions that clearly encodes the direction that was followed during path memorization. Therefore the navigator is driven along a path which is an exact reproduction of the memorized path.

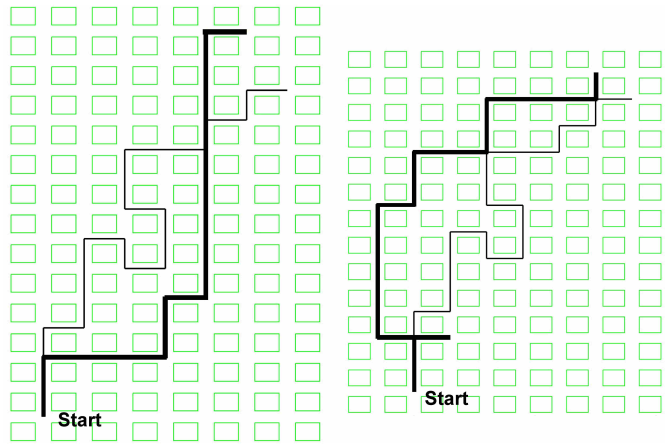

Figure 2: Typical trajectories obtained by the simulated navigator in the Probe Trial condition. In narrow black, the memorized path, in bold black, the simulated reproduced path.

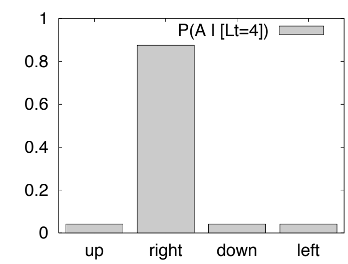

Figure 3: The CPT learned for the term $P(A \mid L_t = 4)$: this is the probability distribution for the fourth intersection of the memorized path (see Fig. 2). At this intersection, the navigator observed twice a movement to the right. This yields a high probability value of going right (0.875, assuming w is set to 0.1).

In the Probe Trial condition, however, landmarks are not available anymore, and $P(A)$ must be computed. In this example, this leads to the probability distribution shown Fig. 4. As the path contains 12 moves in the same direction as the initial orientation (we note this direction "up"), 7 moves to the "right", 1 move "down", and 1 move to the "left", the computation for the term $P(A)$ in the Probe Trial gives the probability distribution shown. At each intersection, we draw at random according to $P(A)$, until 21 displacements have been made: Fig. 2 shows typical trajectories obtained in this manner (bold trajectories).

Discussion

Interpretation of the proposed model

The simulation results illustrate that the proposed computation of the $P(A)$ distribution can be interpreted as a "path integration" component, both in an intuitive sense and in a mathematical sense.

In an intuitive sense, the probability distribution over ac-

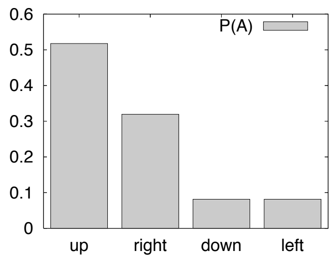

Figure 4: The CPT computed for the term $P(A)$ in the Probe Trial simulation.

tions drives the navigator toward the goal in the correct general direction. Indeed, $P(A)$ as computed can be interpreted as an estimation of the angle from the starting orientation to the goal. For instance, the distribution shown Fig. 4 encodes the knowledge that, to reach the goal, the navigator must mainly go "up" and "right", and it also encodes the relative proportions of these elementary displacements. $P(A)$, seen is this manner, not only encodes an estimate of the bearing of the goal, it also encodes the accuracy or reliability of that estimate, by the spread or uncertainty of the obtained probability distribution.

Let us now recall the mathematical sense of a path integrator. Let $\vec{p}(t)$ be the path, *i.e.* the sequence of elementary displacement vectors at time $t, 0 \leq t \leq T$. The vector $\vec{V}(t)$ representing the global displacement from the initial time 0 to time T is then given by:

$$\vec{V}(t) = \int_0^T \vec{p}(t).$$

In our formulation, time is not continuous, but discretized using events which are the passage at intersections. This explains the use of a discrete summation over intersections L_t instead of a continuous integral over time. Moreover, we assume the elementary displacements are not known deterministically, as in $\vec{p}(t)$, but are encoded using the probability distributions $P(A \mid L_t)$. Therefore, the equation $P(A) \propto \sum_{L_t} P(A \mid L_t)$ can be interpreted as a Bayesian, discrete version of a path integration mechanism.

The simulation shows that our model qualitatively reproduces the patterns of errors made by participants in the Probe Trial. Indeed, in the simulated Probe Trial path reproduction (Fig. 2), we observe that, even though the navigator is driven in the general direction of the goal, the order in which the elementary displacements were performed in the learned path are completely forgotten. This is a direct consequence of the way the probability distribution $P(A)$ is computed. In $\sum_{L_t} P(A \mid L_t)$, the summation can exactly be interpreted as an aggregation of all observed displacements. In other words, the sequencing of displacements, which is present in $P(A \mid L_t)$, is not present anymore in $P(A)$.

The model structure proposes an original hypothesis concerning the interplay between the landmark-based cognitive strategy and the path integration strategy for spatial navigation. Whereas, in the literature, they are commonly pictured as independent mechanisms hierarchically articulated by a main system / back-up system relationship, in our model, there is only one navigation system. When all sensory information are available, this system corresponds to the landmark-based navigation; when some sensory inputs are missing, the same system can operate in a degraded mode, and then exhibits properties of a path integration mechanism.

Relaxing our assumptions: towards experimental predictions and new protocols

We now discuss the relevance of the simplifying assumptions required by our model, which leads us to its possible extensions and the experimental predictions it can provide.

We have assumed, in the model, that all landmarks could be identified with no errors. In a real world navigation scenario, it is of course highly improbable that visual landmarks are never ambiguous. In the COSMOpoliS© experiment, landmarks were unique along the trajectory. However, the study of the way places and intersections are identified is a complete domain of investigation in itself. The goal is to distinguish the intersection identity L_t from the perceived sensory cues at that intersection $P_1, \ldots, P_k$. For instance, landmarks are not the only cues that can be used to identify intersections, as configurations of landmarks could play a role (Mallot & Gillner, 2000), and geometrical configurations of the intersection itself (T-shaped, X-shaped) is probably also encoded (Stankiewicz & Kalia, 2007). In our model, we have assumed that the intersection identity L_t to be readily available; in practice, it could be estimated according to $P(L_t \mid P_1, \ldots, P_k)$. Determining the perceptual components $P_1, \ldots, P_k$ and the structure of this perceptual model is subject of future work.

Another major simplification in our model is the lack of temporal dependency between intersections. Indeed, it is highly probable that pairs $\langle L_t, L_{t+1} \rangle$ of landmarks perceived in sequence, or even higher order sequences $\langle L_t, \ldots, L_{t+m} \rangle$ are used for memorizing paths. Sequences of actions might also serve as large-scale cues for memorizing the paths. This could be incorporated in m-order Markov models of the form $P(L_t, \ldots, L_{t+m}, A_t, \ldots, A_{t+m})$. It might be interesting to use future experimental data in order to estimate m, *i.e.*, the length of the sequences of sensory and motor cues used for path memorizing.

Finally, we wish to discuss the way we generate simulated paths with the model. Indeed, so far, we have assumed the navigation could use probability distributions over actions, and draw at random, at each intersection, directions to follow. In the current simulation, no memory whatsoever is included in this process. In other words, our simulated navigator would not be able to know if it was "unlucky" in its progress, and was deviating away from the memorized orientation to the goal. However, it appears obvious that human navigators would update their estimation of the orientation to the goal as they progress towards it. Mathematically, it would be straightforward to enrich our model to reproduce such a mechanism. Unfortunately, the current experimental data would not enable us to determine the biological plausibility of any such mathematical development.

Conclusion

We have presented a preliminary model of large-scale human navigation in a virtual city. This model successfully qualitatively reproduces patterns of errors that were observed in human participants. In the Landmark condition, where all visual cues are present, both the participants and the simulated navigator accurately reproduce the learned path. In the Probe Trial condition, where the visual cues needed to recognize the current position are missing, both the participants and the simulated navigator are still able to reach the goal, but both do so using variants of the learned paths.

The proposed model is based on Bayesian modeling. A single probability distribution encodes the learned path. It encodes properties of the learned path, and can be used to generate different strategies according to the availability of cues. In the Landmark condition, the probability distribution can be read directly, and the navigator performs as if using a landmark-based navigation strategy. Whereas, in the Probe Trial condition, the probability distribution can be used to generate the best estimate about actions to perform, thanks to Bayesian inference, and the navigation then performs as if using a geometry-based navigation strategy. Having a single model, which is the basis of several navigation strategies, departs from the classical view where each strategy is independently encoded and which requires an arbitrator for hierarchically articulating them.

This could provide novel insights into the cognitive mechanisms involved in human navigation and space representation, and hopefully, could be transferred to biomimetic robotic architectures, where managing hierarchical representations of complex environments is still a challenge.

Acknowledgments

This work has been supported by the BIBA and BACS European projects (FP5-IST-2001-32115, FP6-IST-027140).

References

Berthoz, A., Amorim, M., Glasauer, S., Grasso, R., Takei, Y., & Viaud-Delmon, I. (1999). Dissociation between distance and direction during locomotor navigation. In R. Golledge (Ed.), *Wayfinding behaviour* (pp. 328–348). John Hopkins University Press.

Diard, J., & Bessière, P. (2008). Bayesian maps: probabilistic and hierarchical models for mobile robot navigation. In P. Bessière, C. Laugier, & R. Siegwart (Eds.), *Probabilistic reasoning and decision making in sensory-motor systems* (Vol. 46, pp. 153–176). Springer-Verlag.

Diard, J., Bessière, P., & Mazer, E. (2003). A survey of probabilistic models, using the bayesian programming methodology as a unifying framework. In *The second int'l conf on computational intelligence, robotics and autonomous systems (CIRAS 2003).* Singapore.

Franz, M., & Mallot, H. (2000). Biomimetic robot navigation. *Robotics and Autonomous Systems*, *30*, 133–153.

Hauskrecht, M., Meuleau, N., Kaelbling, L. P., Dean, T., & Boutilier, C. (1998). Hierarchical solution of Markov decision processes using macro-actions. In G. F. Cooper & S. Moral (Eds.), *Proc. of the 14th conf on uncertainty in artificial intelligence (UAI-98)* (pp. 220–229).

Jacobs, L. F., & Schenk, F. (2003). Unpacking the cognitive map: the parallel map theory of hippocampal function. *Psychological Review*, *110*(2), 285–315.

Janzen, G., & Turennout, M. van. (2004). Selective neural representation of objects relevant for navigation. *Nature Neuroscience*(6), 673–677.

Kuipers, B., Modayil, J., Beeson, P., MacMahon, M., & Savelli, F. (2004). Local metrical and global topological maps in the hybrid spatial semantic hierarchy. In *Proc. of the ieee int'l conf on robotics and automation (ICRA04)* (pp. 4845–4851).

Kuipers, B. J. (2000). The spatial semantic hierarchy. *Artificial Intelligence*, *119*(1–2), 191–233.

Lane, T., & Kaelbling, L. P. (2002). Nearly deterministic abstractions of markov decision processes. In *Eighteenth national conf on artificial intelligence (AAAI-2002).*

Mallot, H., & Gillner, S. (2000). Route navigating without place recognition: what is recognised in recognition-triggered responses? *Perception*, *29*, 43–55.

Redish, A. D., & Touretzky, D. S. (1997). Cognitive maps beyond the hippocampus. *Hippocampus*, *7*(1), 15–35.

Stackman, R., & Taube, J. (1997). Firing properties of head direction cells in the rat anterior thalamic nucleus: dependence on vestibular input. *Journal of Neuroscience*, *17*, 4349–4358.

Stankiewicz, B. J., & Kalia, A. A. (2007). Acquisition of structural versus object landmark knowledge. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(2), 378–390.

Taube, J. (1998). Head direction cells and the neurophysiological basis for a sense of direction. *Progress in Neurobiology*, *55*, 225–256.

Thrun, S. (1998). Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, *99*(1), 21–71.

Trullier, O., Wiener, S., Berthoz, A., & Meyer, J.-A. (1997). Biologically-based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, *51*, 483–544.

Victorino, A. C., & Rives, P. (2004). An hybrid representation well-adapted to the exploration of large scale indoors environments. In *Proc. of the ieee int'l conf on robotics and automation (ICRA04)* (pp. 2930–2935).

Voicu, H. (2003). Hierarchical cognitive maps. *Neural Networks*, *16*, 569–576.

Wang, R. F., & Spelke, E. S. (2002). Human spatial representation: insights from animals. *TRENDS in Cognitive Science*, *6*(9), 376–382.