

How *Albot₀* finds its way home:

A novel approach to cognitive mapping using robots

Wai K. Yeap (wai.yeap@aut.ac.nz)

Centre for Artificial Intelligence Research, Auckland University of Technology

Auckland, New Zealand

Abstract

This paper presents a novel approach to cognitive mapping using robots in which robots with different perceptual abilities are used to investigate how they themselves compute and utilize a cognitive map. These robots are thus treated as a new unique species capable of cognitive mapping and are referred to as *Albots*. They are different from other robots which are used to study robot mapping or simulation of theories of cognitive mapping. *Albot₀* is the first such robot created. It is equipped with sonar sensors and an odometer, computes an imprecise cognitive map, and uses it to find its way home. This paper discusses how and why *Albots* are created, and much of what is learned about cognitive mapping from observing how *Albot₀* uses its imprecise map to find its way home.

1. Introduction

A central problem in modeling spatial cognition concerns how a biological agent, be it an ant, a rat or a human, comes to know about its environment. From an initial perception of the environment to its exploration, what is computed and how? Of much interest is the notion of a “map-like” representation, commonly referred to as a cognitive map, and its process, cognitive mapping. Since the seminal work of Tolman (1948), Lynch (1960), and O’Keefe & Nadel (1978) much work has been done by researchers with disparate backgrounds to understand the nature of cognitive mapping and cognitive maps. Among them, and more recently, is the modeling of the cognitive mapping process using robots (for recent discussions, see Jefferies & Yeap, 2008) which is the topic of this paper.

With the development of more powerful and cheaper autonomous mobile robots, researchers are keen to use them for testing ideas about the animal/human cognitive mapping process (for a recent example, see Vasudevan et al., 2007). These studies have been investigated at three different levels: system, neural, and symbolic. At the system level, the robot is created with some physical aspects (usually the sensors) similar to that of an animal or an insect. For example, Lambrinos et al.'s (2000) created a robot that perceives and finds its way like an ant. At the neural level, the focus is on simulating the neural substrate of the brain that is involved in cognitive mapping. Examples of some recent work in this area include Milford and Wyeth (2007), Barrera and Weitzenfeld (2008), and Hafner (2008). At the symbolic level, one is interested in the abstract algorithms for cognitive mapping. Examples of such work include Chown, Kaplan and Kortenkamp (1995), Jefferies, Baker and Weng (2008), and Beeson, Modayil and Kuipers (2010).

One well-known difficulty perceived when using robots to model cognitive mapping is that robotic systems have very different sensors and very limited interpretation of their sensing. This makes it difficult to develop a realistic model, and these models are, consequently, largely ignored by those working with biological systems. Note that while the systems approach somewhat overcomes this problem, their robots are difficult to build and, so far, they have only been able to match the senses of some lower animals and insects. In this paper, I discuss a novel symbolic approach using robots for studying cognitive mapping which does not suffer from this drawback and which I have been developing over a number of years with my students and co-workers (Wong et al., 2005; Wong et al., 2007; Yeap et al., 2008). Its key aspect is to treat each robot, with its own unique set of sensors, as belonging to a new unique species which can wander about autonomously and perform cognitive mapping like a natural species. I will refer to this new species as the *Albot* species and each individual member as *Albot_n*. Thus, the research goal is to create a cognitive mapping process that is

appropriate for each individual robot so that one can study how increasingly complex autonomous systems compute their own cognitive maps and how they use them to generate spatial behavior commonly observed in natural species. The emphasis is on the implementation of cognitive mapping *at the robot level*.

That developing Albots could tell us something useful about cognitive mapping is evident in that much of what we have learned to date about cognitive mapping comes from observing how biological agents, with vastly different sensory and reasoning capabilities (such as bees, ants, rats, fish, and humans), behave and interact in the same physical world that we live in. Each species has evolved a unique solution which is in accordance with the way they perceive and move about in their environment. Yet, while these mechanisms differ and some representations computed are richer than others (comparing say humans and ants), they all make much use of some basic information such as distance traversed, orientations between known locations, landmarks and others. It is evident that not all species will produce and use a spatial map, which raises the question about when and why a cognitive map is computed? If a map is used, how is it derived from the sensors? By developing Albots that utilize similar information for solving similar problems, one can investigate what is possible and what is not. This will provide much needed insights into the computational aspects of cognitive mapping.

The fly in the ointment is that, while one could develop a robot to map its environment, there is no guarantee that the robot produced is an Albot i.e. a robot doing cognitive mapping as opposed to robot mapping. The latter kind of mapping would not tell us much about cognitive mapping. How do we distinguish between the two? Or, more importantly, how do we know that what is implemented is cognitively plausible or relevant? Ideally, the best way is to be able to demonstrate or prove that the same algorithm is being utilized in some biological agent but, unfortunately, we lack any direct mechanism to find out if this is the

case. Worse, despite numerous investigations into the nature of cognitive mapping and cognitive maps, both these notions remain poorly understood (Bennett, 1996). The lack of a clear definition and a useful method of proof make the task of designing Albots extremely challenging.

How then are Albots developed and what would an Albot be like? The remainder of this paper expounds the intricacies of this new approach, presents Albot₀ – the first in the series and discusses what it tells us about cognitive mapping – and, in the conclusion, foreshadows the development of Albot₁, what our next Albot could be like.

2. Albots and related approaches

Albots are created to investigate cognitive mapping at the robot level. As noted earlier, there is a danger that what one produces is just another robot mapping its own world and which has nothing to do with cognitive mapping. It is therefore important to first highlight the key characteristics of this approach that distinguishes it from other uses of robots to do mapping, and which shows that what is developed is cognitively relevant. Furthermore, since all researchers share a common problem but with different goals, this section will also briefly discuss what opportunity there is for cross-fertilization between them.

There are two such approaches that are strongly related to the Albot approach. The first is what I will refer to in this paper as the simulation approach and the second is the general robotics approach. In the simulation approach, robots are used as a medium for implementing ideas/theories about cognitive mapping. The implementation is primarily intended to provide some realism to the theory and to help ground it to the real world. As mentioned earlier, given that the robot sensors are different and that there is no clear-cut criteria to decide if an algorithm is cognitively plausible, many of these researchers pay little attention to the cognitive reality of their algorithms. They also seldom analyze and compare the behavior of

their robots with those of biological ones. Instead, they pay close attention to developing a complete working system that embodies their theory of cognitive mapping. Thus, many, rightly so, claim their work to be cognitively inspired rather than cognitively plausible. Two recent examples of such work include that of Jefferies, Baker and Wong (2008) and Beeson, Modayil and Kuipers (2010). The former implemented an MFIS theory of cognitive mapping (Yeap, 1988; Yeap and Jefferies, 1999) and the latter implemented a hybrid SSH theory (Beeson, 2008).

The simulation approach is thus similar to the Albot approach in that both attempt to implement ideas/theories about cognitive mapping using robots. However, they differ in the way in which they go about implementing these ideas/theories. In the simulation approach, one is concerned with getting these ideas implemented using robots, whereas in the Albot approach, one is concerned with using robots to investigate the truth about the ideas themselves. This does not mean that researchers using the simulation approach are not interested in the truth of the ideas being implemented. Rather, the truth of their ideas is established elsewhere, as part of their theory, and during implementation, they are more concerned with the issue of implementability. In the Albot approach, and as we shall soon see below, the implementation is about finding answers to some important questions about cognitive mapping. It is not just about getting the robot to map its environment successfully. Put another way, using the Albot approach, a useful result is not one in which the mapping problem is simply solved but rather solved in a way in which some insight into the nature of cognitive mapping is gained. The significance of the Albot approach thus lies in its explanatory power rather than its demonstrative power.

The explanatory power of the Albot approach comes as a result of asking detailed questions about how a cognitive mapping process is implemented. To elaborate, consider Beeson, Modayil and Kuipers's (2008) implementation of a local perceptual map (henceforth referred

to as the BMK's approach). When moving about in its environment, an autonomous system, be it a robot or a biological system, needs to integrate successive spatial views to form a local global map of its immediate surroundings. BMK implemented such a map using a method developed by robotics researchers and which has been shown to be an efficient method for computing a large global metric map for robots. This raises an interesting question: if so, why doesn't this local global map become the cognitive map itself? BMK argue against it by pointing out that the algorithm might fail for very large environments and that their theory and others have argued for a topological representation being a more efficient representation for a cognitive map. Nonetheless, the basis of that algorithm is aimed at computing a large accurate global map, and newer versions of it are having increasingly better performance (Durrant-Whyte and Bailey, 2006). BMK's restriction of its use to compute a small global map is an arbitrary decision.

From a cognitive perspective and using the Albot approach, it is important to ask: Could such an algorithm be used in biological systems? Could the biological systems, like BMK, choose to restrict the use of such an algorithm for computing a small global map? Or, could a cognitive map, at least for some lower animals, be a large global map? If not, why not? Most importantly, could we demonstrate, using an Albot, that an alternative and more cognitively plausible algorithm is possible? If so, this might provide clues as to the true nature of this representation. Note that while many researchers working on cognitive mapping, on the one hand, do not believe that a cognitive map is a large global map, most researchers working on spatial updating, on the other hand, consider a cognitive map as a global map with an allocentric frame of reference (e.g Wang and Spelke, 2000; McNamara, Rump and Werner, 2003) There is much confusion in the literature regarding what a cognitive map is and one possible way to resolve this problem is to look at how such a map could be computed. In other words, the choice of an algorithm matters and while we don't know for sure what

algorithm is used in biological systems, we could infer from the implementation what is possible and what is not. If what is inferred does not explain the findings from studies of biological systems, then one needs to develop alternative solutions. If no alternative solution can be found, we need to re-examine our findings from biological systems. Note that one is not looking for an alternative solution which is more efficient but rather one which could explain the findings from studies on biological systems. Asking these questions and finding an answer form the gist of the Albot approach.

In the robotics approach, researchers are concerned with developing efficient methods for the robots to compute a map of the physical environment. While these researchers have little concern over whether their solution is cognitively plausible or not, they nevertheless, like the Albot researchers, are involved with developing solutions for robots to do mapping. Thus, both researchers face many similar problems and while it is unlikely that the robotic solution is cognitively relevant, it is nonetheless an example of how that problem could be solved. This provides a useful starting point for developing Albots and vice versa, the methods developed using Albots would be of interest to robotics researchers. Nature is known to have developed powerful solutions to its problems. Furthermore, as more autonomous humanoids are being developed (example, Michel et al., 2006) and as autonomous robots are performing increasingly more tasks than just mapping their environment, new methods might be needed and nature's solution might be the best.

In summary, the robotics approach is about efficiency, the simulation approach is about implementability, and the Albot approach is about explanatory power. Viewed in this manner, opportunity for cross-fertilization between them becomes clear. For example, on the one hand, using the Albot approach, one could discover how nature solves the mapping problem which robotics researchers could then improve upon. On the other hand, robotics researchers have discovered solutions to the mapping problem that researchers interested in

developing an Albot should review and consider whether they are appropriate and, if not, why not. Similarly, simulation studies of theories of cognitive mapping could highlight issues that require more careful analysis using Albots while insights gained from experiments with Albots help re-shape the theory. Finally, the two different systems approaches for building robots to map their environment, a robotic system versus a cognitively inspired system, could provide different insights into building complex systems with mapping capabilities, and each could learn from one another what works best. Fig. 1 illustrates this three-way exchange of ideas among these different groups.

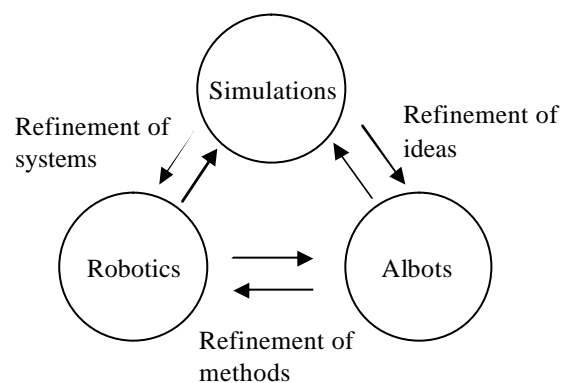


Fig. 1. Cross-fertilization of ideas in mapping research

3. Creating Albots – The basic idea

Albots are artificial creatures (i.e. robots) created for investigating the nature of cognitive mapping. Its advantages are twofold. First, it is a much simpler system (compared to natural species) and one could thereby use it to isolate and focus on the study of some particular issues of cognitive mapping. Second, it forces one to investigate cognitive mapping issues arising from one's perception of the environment. This aspect of cognitive mapping has rarely been studied and it is important to do so. This section looks at the basic idea of creating an Albot – what issues are to be investigated and why.

Past observations about cognitive maps came mainly from behavioral studies of a wide ranging number of biological agents. For general reviews of such work, see Downs and Stea (1973), Gallistel (1990), Poucet (1993) and Golledge (1999). These works have generally suggested a representation with a rich structure consisting of many interesting features. A cognitive map is thus more than a memory of the spatial location of objects perceived. For example, the map is often said to contain information about paths/routes, landmarks, shape/geometry of local space, and hierarchies of place representation. The information in it is known to be imprecise and incomplete and yet biological agents appear to use it effortlessly to return home, to orient to unseen locations and, if needed, to create short-cuts or find new paths. At a more abstract level, significant events that happened in a place are remembered, functions of places are recognized and habitual routines through the environment established.

However, how such knowledge emerges from our perception of, and experience in, the physical environment remains little known. For example, how and when are landmarks formed? What is a place? What is a map-like representation? How does the geometry of each local space visited become part of the path traversed? If our map is incomplete, what is remembered and why? How could one use an incomplete and imprecise map to find one's way home? What kind of representation is needed to support short-cuts and the discovery of new routes?

Albots are created specifically to help find answers to these questions. Their representations of the environment are perceptually driven and each being derived using a very different perceptual system. As such, each Albot could be used to investigate those features of cognitive maps that are computable from the information made explicit in its perceptual system. Albots are thus not created just to demonstrate how cognitive mapping is implemented in an artificial system. For example, consider again the problem of computing a

local global map as discussed in the previous section. Its underlying problems are how different views are integrated to form such a map and how such a map could eventually lead to the construction of a cognitive map that has imprecise and incomplete information. To study this problem, one could create an Albot whose perceptual system uses a laser sensor as its primary sensor (see the discussions of Albot₁ below). The latter is chosen because it makes explicit the key information needed to study this problem i.e. the spatial layout of the surfaces in view. It is not chosen to investigate how cognitive mapping, in general, could be implemented in a robot equipped with laser sensors. The latter would draw the kind of criticisms mentioned earlier (such as the dissimilarity of sensors and the lack of conceptual knowledge).

Another example of an Albot is one created to address an even more fundamental question about cognitive maps: its imprecise nature. Much has been said about a cognitive map being an imprecise representation but what exactly does this mean? Could one still use such a map to move about in one's environment? In robot mapping, researchers avoid using such a map simply because it would be difficult especially when no other useful information (such as objects being recognized) is available. So, how does a biological agent use such a map to solve many of the spatial tasks observed? Some researchers working with biological agents have begun to argue that other non-spatial information is used and that such a map does, in fact, not exist (Benhamou, 1996; Bennett, 1996; Mackintosh, 2002; Foo et al., 2005). As we shall see below, Albot₀ is created to address this question. It uses sonar sensors to create a highly imprecise representation of its environment and the problem of using such a representation to find one's way home is then investigated.

Once a cognitive map is computed using an Albot, it is important to show how the map is used to reproduce the same kind of spatial behavior as observed in biological agents' use of their cognitive map. Without doing so, and as already highlighted in our earlier discussion of

the implementation of the local global map, the map computed might not be a cognitive map. Thus, in addition to conducting experiments showing how each Albot computes its map, it is also important to conduct further experiments to show how each Albot uses it to perform spatial tasks typically performed by biological agents. Some of these tasks include finding one's way home, finding short-cuts, finding novel paths, and having the ability to orient towards hidden goals. Note that these tasks are easily solved using a complete and precise map but, again, it is important to stress that one's cognitive map is incomplete and imprecise. Biological agents' performances in these tasks show limitations. For example, some subjects find it difficult to orient to familiar locations and some fail to find short-cuts. Could such behavior be reproduced? If so, it may provide the much needed insights about cognitive mapping. This is because existing debate on the nature of cognitive mapping has been based upon one's inference as to what a cognitive map is from the behavior observed. Doing so, one could easily mistake what map is actually used by the biological agent in question. With Albots, one could discuss these cognitive mapping issues with precise knowledge of what the process is and what map is being computed.

In summary, when creating an Albot, one needs to address three key questions: (i) what features/issues of a cognitive map are being investigated, (ii) what kinds of algorithms are available for computing the cognitive map, and (iii) what kinds of spatial behavior could be generated? Unlike those who aim to reproduce the algorithms used in some natural species (such as the ants (Lambrinos, et al., 2000)), both the map and the algorithms developed in an Albot need not be exactly the same as those found in some biological agents. Albot is a new unique species and one can expect its algorithms to differ due to the way it perceives and moves about the environment. As long as the questions addressed are those raised in the studies of cognitive mapping and the key information used to compute the map is the same as those utilized by biological agents, then that Albot is said to compute a cognitive map. It is in

this sense that the implementation of an Albot is cognitively relevant. After all, different species could use the same sensors (vision, for example) in different ways and compute a cognitive map based on their own need to survive.

4. Albot₀ – An Example

The “birth” of Albot₀ was an afterthought. Initially, together with my co-workers and students, I was testing my computational theory of cognitive maps (Yeap, 2007; Yeap, 1988; Yeap and Jefferies, 1999) using a mobile robot. One central idea of my theory was that the cognitive mapping process begins with the learning of a network of local spaces (or, what I referred to as a network of absolute space representations or ASRs). The question we asked was: how do we get a robot to compute a similar network of ASRs as its cognitive map? We note that robotics researchers have long been developing algorithms for partitioning the environment into smaller spaces (for a review, see Yeap and Jefferies, 2000). However, their research has turned this problem into what they called SLAM, simultaneous localization and mapping (for a recent review, see Thrun, 2008). That is, they are primarily concerned with the robot not getting lost and in particular, when it re-visits a familiar part of the environment, it should not fail to recognize it. This is achieved by knowing its position in the map. Consequently, much of their effort is focused on correcting errors in sensor measurements.

However, few biological agents, if any, have been reported in the cognitive mapping literature that demonstrates the use of an accurate metric map. Insects, notably ants, come close but they only tally the distance traversed to do path integration and not to build an accurate metric map. Nature is telling us we don’t need a metric map. Thus, our challenge was to compute a network of ASRs from sensor measurements as they are (i.e. without error correction). What kind of a map could we really compute and could it be useful? To make it even more challenging, we equipped our robot with only sonar sensors and an odometer.

Both are known to be highly inaccurate and this will ensure that whatever map is created will be fuzzy and one could not read off from the map the precise location of where one is. Furthermore, since the map created will be so impoverished, the solution for finding one's way home, if there is one, must make much use of the spatial information afforded in it. We successfully created an algorithm for the robot to compute a network of fuzzy ASRs and an algorithm for it to use the network to find its way home. Most interestingly, the algorithms developed captured some interesting characteristics of cognitive mapping as observed in biological agents. For example, the map computed is not only fuzzy but a different network of ASRs is computed depending on how the robot is travelling, whether in the outward direction or in the homeward direction. The returning home algorithm made much use of two key pieces of information found in cognitive maps, namely, exit and distance information derived from the spatial extent of each local space visited.

The resulting robot could be thought of as a kind of a rat and just like, say, a mole rat, it is different from an albino rat or a household rat because its sensors are different. It is not ready to perform the amazing feat of the laboratory rat in a maze but it does demonstrate that a spatial map computed from an almost featureless environment could still be useful to solve the very basic cognitive mapping task, namely returning home. With hindsight, the question we asked and the experiments conducted represent the gist of what I now described as the Albot approach to cognitive mapping. This robot thus becomes the first robot in the Albot series and is now described here as Albot₀. More technical details on Albot₀, can be found in Schmidt et al. (2007), Wong et al. (2007) and Wong (2009).

4.1 The robot

The robot used is a Pioneer 3DX mobile robot (from MobileRobots Inc.). It is equipped with a ring of eight sonar sensors and an odometer. With sonar, not much is perceived when

the robot is standing still. Thus, to create a “view”, Albot₀ has to move from one part of the environment to another. Albot₀ is allowed to wander freely on its own, stops whenever it encounters an obstacle, turns and moves again, maintaining a general forward direction. Sonar data collected from each start-stop movement constitutes its view of the environment. Fig. 2 shows Albot₀’s environment and a typical path through it, and the sonar points obtained. These points are grouped together to form lines denoting surfaces in the environment.




Fig. 2. Albot’s environment and a typical path (as indicated by the solid arrow line) through (left) and the sonar data obtained (right). Note how messy the data points are. The paths traversed in most experiments are usually longer than 50m.

4.2 The map

For each view, Albot₀ computes the rough shape of the local space around it. This is achieved by connecting surfaces perceived on each side of the path to form a boundary for the local space. The perceived surfaces are viewed at four different resolutions. Surfaces of the highest resolution (i.e. the largest surfaces) are grouped together first to form an initial

boundary. For each side, if the total length of those collected surfaces occupied more than 70% of the path traversed, then the boundary is said to be indicative of the boundary of the current local space. Otherwise, surfaces of lower resolutions are added iteratively to form the required boundary. Note that the 70% criterion is an arbitrary one and is meant to emphasize that all we need is a rough description of the boundary

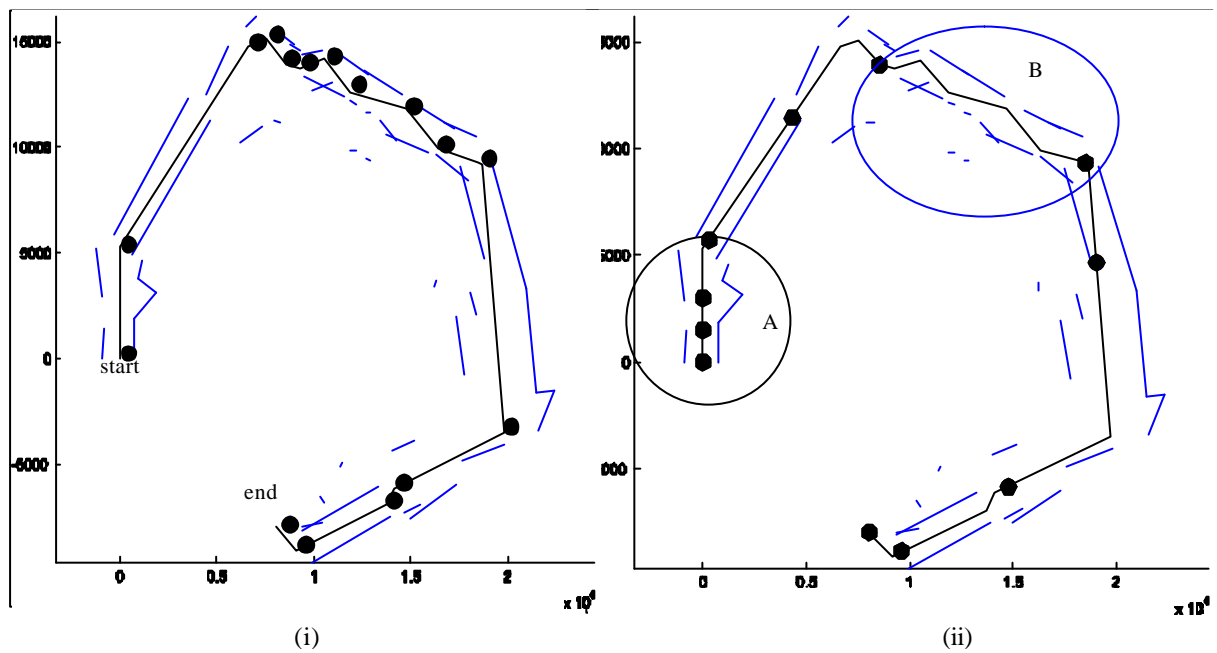


Fig. 3. (i) A network of local spaces computed from each view in an outward journey, and (ii) which is then turned into a network of ASRs. (0,0) indicates the home position of the robot. The lines connecting the filled circles indicate Albot₀'s path and the space in between denote the local space or ASR computed respectively. The remaining solid lines indicate the rough shape of each local space or ASR computed.

The local spaces derived from each subsequent views of the environment need to be combined to form a view-independent description of the current local space in which the robot is enclosed. Such a representation is what I called an ASR. Thus, as Albot₀ moves forward, these local spaces (from each view) might either be split to form different ASRs or merged to form a single ASR. Fig. 3 shows examples of a split and a merge situation and the network of ASRs computed for the path shown in Fig. 2. In the region marked (A) where Albot₀ enters an open space from a corridor and leaves straight into another corridor, the local

space computed was split into three different ASRs. In the region marked (B), five local spaces resulting from zigzag movements down a corridor were combined to form a single ASR.

4.3 The journey home

To find its way home using the map created, Albot₀ could not rely on the shape of the ASRs computed. What is more, it also could not rely on counting the number of ASRs computed. It turns out that in the homeward journey one could possibly produce more, or less, ASRs, a consequence of having a different perspective of the path traversed. This, I believe, makes Albot₀ behavior more cognitively interesting. Taking a cue from animal behavior, we sought and developed an algorithm for returning home that focused on the distance traversed and the orientation between adjacent ASRs. The latter is not expected to be very useful in this environment (i.e. no sharp turns) but nonetheless, boosts Albot₀'s belief of where it is in its homeward journey. Given the map computed, we found that one effective measure of distance is the straight line distance traversed in each ASR, from its entrance to its exit. Measured in this way, Albot₀ remembers the total distance traversed in the outward journey and as it returns home, it travels the same distance, calculated from the ASRs computed in the homeward journey.

For example, consider the two maps computed in Fig. 4. Each pair of dots indicates the entrance and exit of an ASR. The distance as measured from the zig-zag movements of the robot in between these dots would not be reliable to use for returning home. Instead, what is used is the sum of the straight line distance as measured between the dots themselves. The basic strategy then is to compare the distance d traveled when returning home, measured in ASR lengths taken from the intermediate map (the map on the right of Fig. 4) computed on the return journey, to the ASR lengths taken from the map computed in the outward journey (the map on the left of Fig. 4). It turns out that this algorithm is effective and Albot₀ manages

to find its way home in 80% of more than 20 trials in the corridor of the environment as shown in Fig. 2 (Wong, 2009, p.102). In another experiment involving more than 10 trials, Albot₀ went through the large empty room in the middle of the environment and was able to find its way home with a success rate of 70% (Wong, 2009, p.114). An example of one such trial of going through the middle room is as shown in Fig. 4.

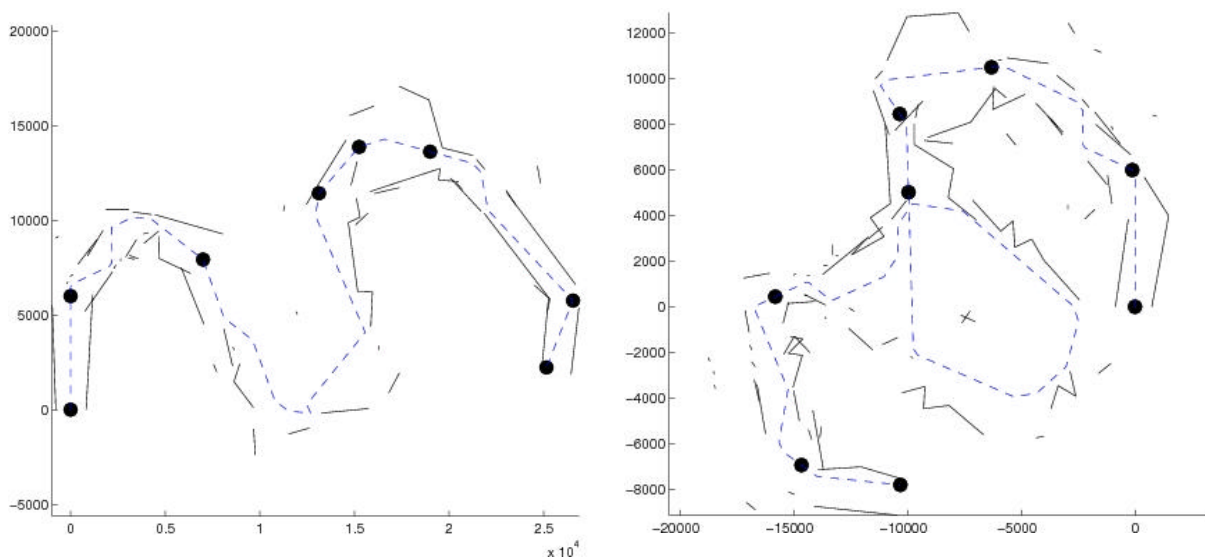


Fig. 4 Albot₀'s outward (left) and homeward (right) journey involving going through the large middle room in the environment of Fig. 1 and the respective network of ASRs computed. Albot₀ successfully found its way home in this experiment.

5. Discussion

Albot₀ was created to study the single most important aspect of cognitive mapping, namely its imprecise spatial nature. Having a cognitive map implies having a broader sense of where things are in one's environment and in particular, one could use it to find short-cuts if needed (Tolman, 1948; O'Keefe and Nadel, 1978). Yet, the precise nature of such a representation and how it is derived from one's direct experience of the environment has been unclear (Yeap and Jefferies, 2000; Mackintosh, 2002). Furthermore, in recent years, questions have been raised as to whether animals, and even humans, have a cognitive map (Benhamou, 1996; Bennett, 1996; Mackintosh, 2002; Foo et al., 2005).

In creating Albot₀, it was shown how a spatial map is computed for a simple robot equipped with sonar sensors and how Albot₀ becomes a species which can roam in a corridor-like environment and find its own way home. Two basic steps are found to be important for Albot₀ to create its spatial map. First, “larger” surfaces perceived in each view are remembered and are used to construct a rough shape of the current local space visited. Second, a “split and merge” algorithm is used to turn these local spaces into a view-independent description. Due to its poor sensors, the map computed by Albot₀ is extremely impoverished; the map neither has accurate position information of the surfaces in it nor an adequate description of the shape of each local space for later recognition. Yet, it was shown that the map still affords useful distance information, as measured between entries and exits of local spaces, which enables Albot₀ to find its way home. It is interesting to note that while researchers working on cognitive mapping put much emphasis on the use of distance information, their focus was primarily on the biological agents’ knowledge of distance between some known points in space (such as the location of a food source or a hidden platform and/or between known objects/landmarks) and/or how they use distance information to do path integration. Few have thought of it as information implicitly available in a spatial map or considered that distance between the entry and exit of a local space could play such an important role.

The fact that Albot₀’s map which captures only a rough spatial extent of the environment experienced is found to be useful for returning home is a strong indication of the importance of the spatiality afforded in a cognitive map. For biological agents with sensors that could capture better the spatial extent of their environment (as opposed to Albot₀’s sensors), one could expect much more powerful means of utilizing such information. This finding is important in light of recent doubts about whether a cognitive map is a useful term to have and whether cognitive mapping exists in animals (Benhamou, 1996; Bennett, 1996; Mackintosh,

2002). One powerful alternative often suggested by these critics is the use of path integration, possibly together with recognition of some landmarks and the use of other non-spatial reasoning (see below). While this alternative is always a possibility, and is useful to explain animals' behavior in mazes, our results show one cannot rule out the use of a spatial map even in such situations. Albot₀'s map is computed as a result of exploring similar environments i.e. ones with little distinguishable features.

Furthermore, critics of cognitive mapping only point out that existing experiments do not prove that animals have a cognitive map. They have not produced evidence against the use of such a map. For example, typical biological subjects in many of these experiments are trained in one or more fixed routes to the goals and their spatial memory is later tested in terms of their ability to locate one or more of these goals. Much of what happens in the journey is usually ignored (unless they are trained to take note of some landmarks/objects along the path). While the task performed is spatial, some researchers have, rightly so, argued that these experiments do not rule out the subjects' use of other non-spatial methods or spatial methods that do not require a map to perform the task. For instance, in deciding where to search for food next in a radial arm maze, the rats may have considered only which alley they have not visited before rather than be using a spatial map of places marked visited and not visited (Brown, 1992; Brown and Drew, 1998). Similarly, when searching for a hidden platform in the popular Morris water maze, the rats may have figured out where the platform is from their perception of the maze as opposed to from their spatial memory of the whole environment experienced (Benhamou, 1996). Even for the acid test for cognitive mapping, the ability to find short-cuts could be due to path integration rather than spatial reasoning involving a cognitive map (Bennet, 1996).

These critical reviews and others point towards a serious lack of consideration in many of these experiments regarding what the underlying representation of a cognitive map might be

and in particular, the very spatial nature of such a map. Without these, the results of these experiments are open to (mis-)interpretations. It is important that future experiments must focus on testing the use of information that could be shown to derive from a spatial map of the environment. Nonetheless, there are a few experiments in the past that placed emphasis on the spatial quality of the map and they provide some fascinating insights. A classic example is Cheng's (1986) experiment with rats locating food in a simple rectangular enclosure and the subsequent discovery that rats do pay attention to the geometrical shape of the room. A more recent example, and equally fascinating, is Menzel et al.'s (2000) experiment with honeybees that does not require them to learn a fixed route to a feeding station and the subsequent discovery that honeybees do have a spatial memory of their surroundings. These experiments are not easy to design and it is hoped that future Albots might help to provide insights and ideas for experimentation.

When wandering in a corridor-like environment, the map serves Albot₀ well; it enables Albot₀ to find its way home. However, it does not allow Albot₀ to take short-cuts and it does not support Albot₀ wandering into more complex environments. Many earlier comments on our work have pointed out that the success of our algorithm is very much due to the peculiar nature of our test environment – Albot₀ rarely wanders out of the corridor. Consequently, many (and especially robotics and AI researchers) have pointed out that our algorithm is limited and would fail in more complex environments. There is little doubt that Albot₀ is basically moving through a corridor and it could not “survive” beyond a corridor-like environment but then, neither could, say, an ant be able to find its way home if it is being purposely displaced in its journey nor could a blind mole rat survive if it is taken out of its complex underground burrows. It is important to stress that the development of Albot₀ is neither about finding an efficient algorithm for it to wander in highly complex environments nor to compare its performance with other robots in the hope of finding the best algorithm for

robots. If the latter, then it is appropriate to ask how well Albot₀ could find its way home or why it fails in some of its trials. It is not and therefore these questions are inappropriate here (at least for now). Rather, one observes Albot₀'s cognitive mapping behavior in its own environment and discusses the insights gained about cognitive mapping in general. Thus, like a fish that evolves to live in water and a desert ant in a desert, Albot₀ is created to "live" only in a corridor environment. Furthermore, just like biological agents, Albot₀ can sometimes get lost although, unlike biological agents, Albot₀ does not have other means to help it to relocate home when it is lost. This perhaps could be one of the next challenges for designing some future Albots.

6. Conclusion and Future Work

A novel robotic approach for investigating cognitive mapping is presented whereby robots are created with their own built-in cognitive mapping. Such robots are referred to as Albots. Albot₀ is the first such robot and it exhibited the following characteristics of cognitive mapping: (i) having an imprecise map, (ii) the way the map is computed is dependent on the direction of travel, and (iii) much use is made of the distance information between exits and distance to find its way home. What Albot₀ tells us about cognitive mapping is that one should not yet rule out the use of a spatial map in animals, and that future experiments designed to test the presence of such a map using biological agents must demonstrate that these subjects do make much use of the spatial information afforded in such a map. For instance, the distance between entrance and exits of each local space could be very useful.

Albot₀ is equipped with sonar sensors to investigate the problem of how a fuzzy view of the world could still give rise to a useful cognitive map. By replacing the sonar sensors with other kinds of sensors, one could create different Albots to investigate other issues related to cognitive mapping. Future Albots in this series could include one equipped with lasers to

investigate the problems of computing an appropriate local global map for cognitive mapping (as discussed earlier), and one equipped with vision to investigate the effect of objects/landmarks recognition in cognitive mapping. My students and I have already begun work on the former problem and we have assigned the task to Albot₁. Briefly, we note that SLAM produces a class of solutions which focuses on integrating successive views to form a global representation. This solution might be called viewer-centered in the sense that the representation computed is an integration of one's *views* of the environment. An alternative solution could be described as object-centered in the sense that the representation computed is via recognizing what is in the environment (as opposed to knowing where the next view is). The latter is possible by assuming that the environment one lives in is relatively stable and there is much overlap between views to enable us to keep track of how the environment is expanding (via recognizing familiar objects between views). A new algorithm based upon this idea has been implemented to compute a local global map for Albot₁. We are currently conducting spatial experiments with Albot₁ and will report our work in due course.

Acknowledgements

My thanks to all my past collaborators on this project: Dr Chee Kit Wong (past PhD student), Dr Jochen Schmidt (past post-doctoral fellow), and Professor Sapiyan who spent a sabbatical year at the Centre. I would also like to thank three anonymous reviewers and Dr Gunzelmann who provided some very useful comments for improving the presentation of this paper.

References

Barrera, A., & Weitzenfeld, A. (2008). Biologically-inspired robot spatial cognition based on rat neurophysiological studies. *Autonomous Robot*, 25, 147-169.

- Beeson, P. (2008). Creating and utilizing symbolic representations of spatial knowledge using mobile robots. *PhD Thesis*. The University of Texas at Austin, USA.
- Beeson, P., Modayil, J., & Kuipers, B. (2010). Factoring the mapping problem: Mobile robot map-building in the hybrid spatial semantic hierarchy. *International Journal of Robotics Research, 29*, 428-459.
- Benhamou, S. (1996). No evidence for cognitive mapping in rats. *Animal Behavior, 52*, 201-212.
- Bennett, A.T.D. (1996). Do animals have cognitive maps? *The Journal of Experimental Biology, 199*, 219-224.
- Brown, M. (1992). Does a cognitive map guide choices in the radial-arm maze? *Journal of Experimental Psychology: Animal Behavior Processes, 18*, 56-66.
- Brown, M., & Drew, M.R. (1998). Exposure to spatial cues facilitates visual discrimination but not spatial guidance. *Learning and Motivation, 29*, 367-382.
- Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition, 23*, 149-178.
- Chown, E., Kaplan, S., & Kortenkamp, D. (1995). Prototypes, location, and associative networks (PLAN): Towards a unified theory of cognitive mapping. *Cognitive Science, 19*, 1-51.
- Downs, R.M., & Stea, D. (1973). *Image and environment: Cognitive mapping and spatial behavior*. Chicago, IL: Aldine Publishing.
- Durrant-Whyte, H. and Bailey, T. (2006) Simultaneous localization and mapping: Part 1. *IEEE Robotics and Automation Magazine, June*, 99-108.
- Foo, P., Warren, W.H., Duchon, A., & Tarr, M.J. (2005). Do humans integrate routes into a cognitive map? Map- versus landmark-based navigation of novel shortcuts. *J. Experimental Psychology: Learning, Memory, and Cognition, 31*, 195-215.

- Gallistel, C.R. (1990). *The organisation of learning*. Cambridge, MA: MIT Press.
- Golledge, R.G. (1999). *Wayfinding behavior: Cognitive mapping and other spatial processes*. Baltimore: John Hopkins University Press.
- Hafner, V.V. (2008). Robots as tools for modeling navigation skills – A neural cognitive map approach. In M.E. Jefferies & W.K. Yeap (Eds.), *Robotics and cognitive approaches to spatial mapping* (pp. 315-324). Berlin, Heidelberg: Springer Tracts in Advanced Robotics.
- Jefferies, M. E., Baker, J. T., & Weng, W. (2008). Robot cognitive mapping – A role for a global metric map in a cognitive mapping process. In M.E. Jefferies & W.K. Yeap (Eds.), *Robotics and cognitive approaches to spatial mapping* (pp. 664-672). Berlin, Heidelberg: Springer Tracts in Advanced Robotics.
- Jefferies, M.E., & Yeap, W.K. (2008). *Robotics and cognitive approaches to spatial mapping*. Berlin, Heidelberg: Springer Tracts in Advanced Robotics.
- Lambrinos, D., Moller, R., Labhart T., Pfeifer, R., & Wehner, R. (2000). A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems*, 30, 39-64.
- Lynch, K. (1960). *The image of the city*. Cambridge, MA: MIT Press
- Mackintosh, N.J. (2002). Do not ask whether they have a cognitive map, but how they find their way about. *Psicologica*, 23, 165-185.
- McNamara, T.P., Rump, B., & Werner, S. (2003). Egocentric and geocentric frames of reference in memory of large-scale space. *Psychonomic Bulletin & Review*, 10, 589-595.
- Menzel, R., Brandt, R., Gumbert, A., Komischke, B., & Kunze, J. (2000). Two spatial memories for honeybee navigation. *Proceedings of the Royal Society of London, B.*, 267, 961-968.

- Michel, P., Chestnutt, J., Kagami, S., Nishiwaki, K., Kuffner, J., & Kanade, T. (2006). Online environment reconstruction for biped navigation. *Proceedings of the 2006 IEEE International Conference on Robotics and Automation* (pp. 3089-3094). IEEE Press.
- Milford, M., & Wyeth, G. (2007). Spatial mapping and map exploration: A bio-inspired engineering perspectives. *Proceedings of the 8th International Joint Conference on Spatial Information Theory (COSIT)* (pp. 203-221). Springer Verlag Lecture Notes in Computer Science.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford, England: Oxford University Press.
- Poucet, B. (1993). Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review*, 100, 163-182.
- Schmidt, J., Wong, C.K., & Yeap, W.K. (2007). Spatial information extraction for cognitive mapping with a mobile robot. *Proceedings of the 8th International Joint Conference on Spatial Information Theory (COSIT)* (pp. 186-202). Springer Verlag Lecture Notes in Computer Science.
- Thrun, S. (2008). Simultaneous localization and mapping. In M.E. Jefferies & W.K. Yeap (Eds.), *Robotics and cognitive approaches to spatial mapping* (pp. 13-41). Springer Tracts in Advanced Robotics.
- Tolman, E.C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55, 189-208.
- Vasudevan, S., Gachter, S., Nguyen, V. & Siegwart, R. (2007). Cognitive maps for mobile robots – an object based approach. *Robotics and Autonomous Systems*, 55, 359-371.
- Wang, R.F. & Spelke, E.S. (2000). Updating egocentric representations in human navigation. *Cognition*, 77, 215-250.

- Wong, C.K., Yeap W.K., & Sapiyan M. (2005). A mobile robot that maps naively but plans intelligently. *Proceedings of the Artificial Intelligence and Application Conference* (pp. 562 – 567).
- Wong, C.K., Schmidt, J., & Yeap, W.K. (2007). Using a mobile robot for cognitive mapping. *Proceedings of the 20th International Joint Conference on Artificial Intelligence* (pp. 2243-2248), AAAI Press.
- Wong, C.K. (2009). Cognitively inspired mapping on a mobile robot. *PhD Thesis*. Auckland University of Technology, New Zealand.
- Yeap, W.K. (1988). Towards a computational theory of cognitive maps. *Artificial Intelligence*, 34, 297-360.
- Yeap, W.K., & Jefferies, M.E (1999). Computing a representation of the local environment. *Artificial Intelligence*, 107, 265-301.
- Yeap, W.K., & Jefferies, M.E (2000). On early cognitive mapping. *Spatial Cognition and Computation*, 2, 85-116.
- Yeap, W.K. (2007). From spatial perception to cognitive maps: How is the flow of information controlled? In H. Schultheis, T. Barkowsky, B. Kuipers & B. Hommel, (Eds.), *Control mechanisms for spatial knowledge processing in cognitive/intelligent systems* (pp. 59-61). AAAI Press.
- Yeap, W.K., Wong, C.K., & Schmidt, J. (2008). Using a mobile robot to test a theory of cognitive mapping. In M.E. Jefferies & W.K. Yeap (Eds.), *Robotics and cognitive approaches to spatial mapping* (pp. 281-296). Berlin, Heidelberg: Springer Tracts in Advanced Robotics.