Initial Experiments with a Mobile Robot on Cognitive Mapping

Wai K. Yeap, Chee K. Wong, and Jochen Schmidt

Robotics Research Group
Institute for IT Research
Auckland University of Technology
Auckland, New Zealand
{wai.yeap,chee.wong,jochen.schmidt}@aut.ac.nz

Abstract. This paper shows how a mobile robot equipped with sonar sensors and an odometer is used to test ideas about cognitive mapping. The robot first explores an office environment and computes a "cognitive map" which is a network of ASRs [1]. The robot generates two networks, one for the outward journey and the other for the journey home. It is shown that both networks are different. The two networks, however, are not merged to form a single network. Instead, the robot attempts to use distance information implicit in the shape of each ASR to find its way home. At random positions in the homeward journey, the robot calculates its orientation towards home. The robot's performances for both problems are evaluated and found to be surprisingly accurate.

1 Introduction

This symposium posed an interesting question: how could robotics researchers develop autonomous mobile robots that understand what they are doing and have self-awareness? An answer to such a question is not easy; many would say impossible. The field of Artificial Intelligence (AI) begun with a quest for an answer to the question: How could a program be intelligent? Fifty years on, we readily embodied our intelligent software in the form of an autonomous mobile robot. Rightly so, a challenging question for the next fifty years would be: how could such a robot be intelligent? And, dare we be bold and ask: could such a robot behave with original intent [2]?

It is still a long way before we even begin to understand how we might approach developing such a robot. However, we could now begin conducting numerous experiments on how a robot might, and could, behave like cognitive agents, be they humans or animals. In this paper, we describe one such work - how to use a mobile robot to test a theory of cognitive mapping. Since Tolman [3] suggested that animals (including humans) create a representation(s) of the environment in their minds and referred to it as a "cognitive map", many psychological experiments have been conducted to study the nature of cognitive maps (for a recent review, see [4]). Some of the important characteristics of cognitive maps highlighted by these studies include:

1. distorted information about distances and directions
2. landmarks, places, paths

3. hierarchical organization
4. multiple frames of reference

Many models of cognitive maps have also been proposed and one idea appears to be most prominent, namely that the map begins with some form of a network of "place representations". Examples of models that favor this idea include that of Chown, Kaplan and Kortenkamp [5], Kuipers [6], Poucet [7], and Yeap and Jefferies [1].

However, few of these models, if any, have been developed and tested using a mobile robot to the extent of showing explicitly that such a network could be used to explain and/or account for cognitive mapping behavior. It should be pointed out that it is not the case that these researchers fail/neglect to test their models. Rather, much of the effort to date has been focused on computing a network to demonstrate its use for successful navigation (by a robot). More recent examples of such works include Kortenkamp [8], Jefferies, Cree, Mayor and Baker [9], and Kuipers, Modayil, Beeson, MacMahon and Savelli [10].

Without doubt, if the model was tested using a mobile robot, there is a natural tendency to use the map computed to immediately perform successful navigation tasks. The perfect memory of a robot together with the use of some reasonably accurate sensors/sensing algorithms would make it attractive to do so. However, and consequently, in many instances, the researchers ended up solving a robot mapping problem rather than a cognitive mapping problem. Little effort has been spent using their results to seriously explain and/or account for known cognitive mapping behavior.

It is interesting to note a parallel here between early AI researchers interested in computer vision and AI researchers now interested in using a robot to do cognitive mapping. Researchers then were rightly concerned that computer vision is not the same as image processing (see for example, [11]). Researchers now should also pay attention to the more complex nature of cognitive mapping as opposed to robot mapping.

To understand cognitive mapping using a mobile robot, the map created by the robot should exhibit some of the characteristics of cognitive maps mentioned above. The map should then be used by the robot to find its way home. In conducting the latter experiment, some parallels should be drawn between cognitive mapping in animals and in robot.

This paper is about how to use a robot to test theories of cognitive mapping. We first observed that a key feature of cognitive maps is that the information in them is distorted. A robot equipped with sonar sensors and an odometer would often receive distorted information about its environment. The question is: How does a robot compute a "distorted cognitive map" from such information? How is the map represented? We also observed that for lower animals, much is made use of the distance and direction information implicit in such a map to find their way home. Could our robot find its way home using its distorted map? Could it also be utilizing the distance and direction information implicit in these maps? If so, how does it do it? The rest of this paper describes our experiments with such a robot and concludes with a discussion of some insights gained from the experiments.

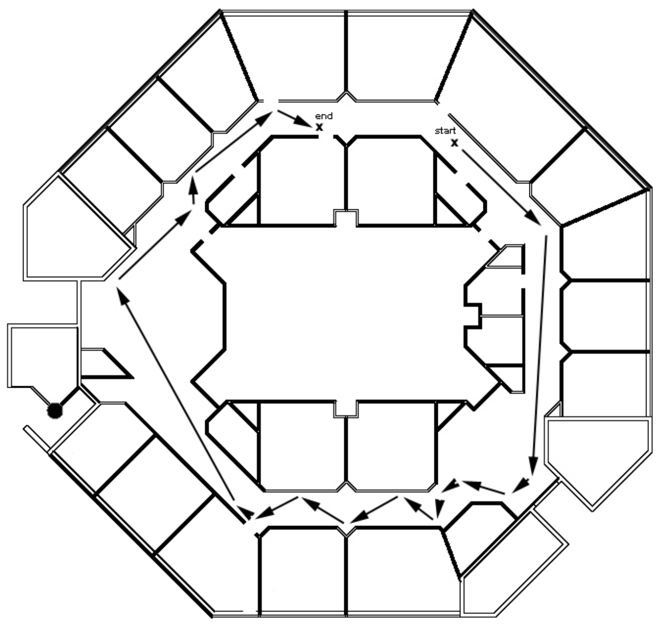

Fig. 1. The environment and the path traversed. The total distance traveled is about 70m.

2 The Robot and Its Cognitive Map

The robot we use is a Pioneer 2 robot from ActivMedia and it came with a ring of 8 sonar sensors. The robot is positioned somewhere in the corridor in an office environment and is allowed to explore the environment until it is told to stop. No modification of the environment is done. That is, things that already existed in the environment (such as rubbish bins, flower pots, cabinets, etc.) remain there and doors leading into offices are close or open depending on the time of the experiment.

The environment used and one of the paths the robot took is as shown in Fig. 1. It does not use a wall-following procedure to navigate. It simply moves forward until it could not and then it ”looks” for an empty space to move forward again. “Looking” is done using all the eight sensors but information about the environment is sensed via the two side sensors. The exploration algorithm used is described as follows:

1. move in a “straight” line and collect sonar data from the sides;
2. stop when an obstacle is encountered; and
3. turn away from the obstacle and continue the mapping process.

Given the above experimental set-up, how does our robot compute a “cognitive map” of its environment? Following Yeap and Jefferies’s theory of cognitive mapping [1], our robot computes a network of local spaces visited (see also [12]). Each local space is referred to, in the theory, as an absolute space representation (or ASR, in short). The details of our new algorithm for computing ASRs for the experiments conducted here can be found in [13–15]. Briefly, the key ideas underlying our new algorithm are as follows:

1. ASRs are computed for each path traversed — a path is a single continuous movement of the robot through the environment (i.e. without any stopping or turning);
2. The important exits found in a path are the exits at both ends of it (i.e. given the poor sensing, one cannot trust the side exits detected). This means that the required ASR for a path is the bounded region for the path;
3. To compute the bounded region, preference is given to using the large surfaces as opposed to the smaller ones, as small surfaces are more likely to be caused by incorrect sonar readings. The algorithm thus uses all the larger surfaces, say, greater than 700mm in length, to compute a boundary. If the resulting boundary is greater than, say, 70% of the distance traveled, then that is an acceptable boundary for the current ASR. If not, more of the small surfaces are added until a reasonable sized boundary is obtained. These figures were found experimentally and work well in practice. Changing their values influences the sizes of the surfaces in the final map.
4. An ASR computed for a path represents an ASR computed for a single view of the robot. The next step is to merge or split ASRs obtained from individual paths into ASRs for the environment experienced. The final ASRs are then connected as a network of ASRs.

Figure 2 shows the final ASRs computed for the journey as shown in Fig. 1. The start and end point of an ASR are marked with a dark circle. The surfaces in between indicate the rough shape of the ASR computed.

3 Theoretical Considerations

In the first outward journey, the robot creates a network of ASRs. What happens in the homeward journey? Obviously, using the same cognitive mapping algorithm, it will produce a network of ASRs for the homeward journey. Figure 3 shows the network of ASRs generated in the homeward journey. Note that 9 ASRs are computed in this instance instead of the 10 ASRs in the outward journey. This difference suggests that not only the shape of the individual ASR could change during the different journeys but also the way in which the path is partitioned could differ as well. Note too that the sixth ASR computed in the outward journey is now perceived as 3 ASRs in the homeward journey.

One strong feature of Yeap and Jefferies's [1] theory of cognitive maps is not simply that the cognitive map is a network of ASRs. Rather, the notion of an ASR is an important representation on its own. It is computed to represent the current bounded local environment that the autonomous system is in (be it an animal or a robot). Just like the information afforded in the current view which is used primarily for solving problems pertaining to the view, an ASR is a representation for solving problems pertaining to the current local environment. For example, it is for telling us, among other things, what lies behind us or what to expect when we turn.

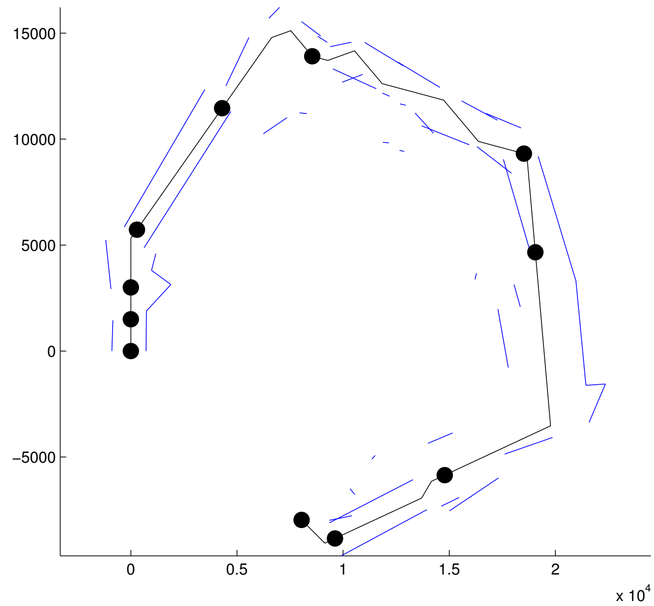

Fig. 2. ASRs computed for the journey as shown in Fig. 1. An ASR is between two adjacent dots and surfaces located to the left and to the right of the path inside the ASR are its boundary surfaces. The path is the solid line connecting the dots. (0,0) indicates the starting position of the robot.

How a cognitive map could emerge from sequences of ASRs experienced in different journeys through the environment is still very much a mystery. It is more than just connecting the different ASRs together as a network of ASRs. Building a cognitive map takes a longer time (than, say, the time spent visiting a local environment). One has to at least ensure that one is settling in a place before remembering the place well. The environment needs to be thoroughly investigated to ensure correct interpretations of the significance of each location. Such understanding will affect the way we perceive the environment and subsequently our construction of a cognitive map representation for it.

If we move away from conventional wisdom whereby one would immediately try to combine ASRs computed in the homeward journey with those computed in the outward journey, what is left? One possible solution is to make use of a key piece of information implicitly available in ASRs, namely the approximate shape of an ASR tells us a good estimate of the approximate distance traveled. This distance estimate is better than the summing up of the zigzag movements of the robot. We implement an algorithm that exploits the use of this information to navigate home.

4 Implementations and Results

In the following, we will present the implementation and results for using the cognitive map generated during the outward journey to obtain the estimates for the current posi-

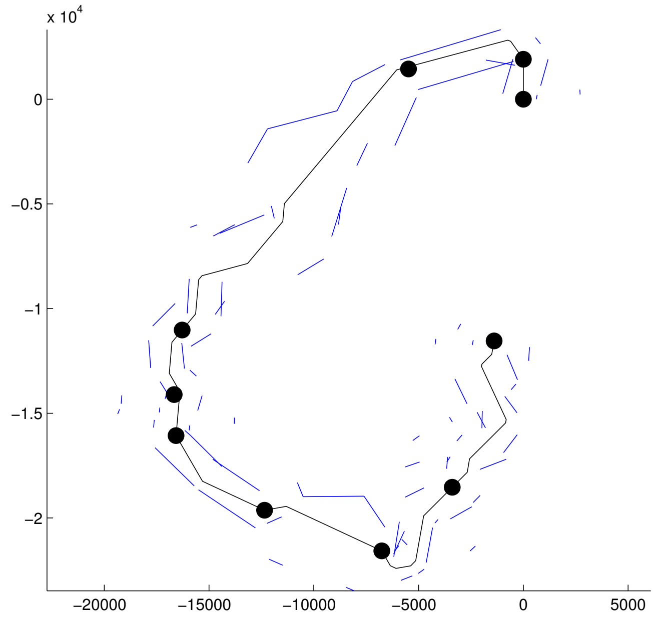

Fig. 3. ASRs computed in the homeward journey.

tion of the robot as well as for its orientation. A total of ten experiments were conducted, results for two of them are described in more detail.

4.1 Going Home

The algorithm for returning home is described as follows:

1. Compute ASRs (up to the current position) in the homeward journey.
2. Measure the length of each ASR computed (as opposed to the actual distance the robot traveled).
3. Map the ASR-distance traveled onto the network of ASRs computed in the outward journey.

The last step provides an estimate of where the robot thinks it is in the actual environment. The robot stops when it believes it has reached home.

The robot computes a cognitive map in the outward journey as shown in Fig. 2. It then makes an attempt to return home, generating new maps and ASRs. Two such maps computed after the robot believes that it reached home are shown in Figures 3 and 4.

We measured the distance between the robot's final position and the actual home position in the real world. For the two experiments presented in the figures, the robot was 1.5m short of the home position for the experiment corresponding to Fig. 3 and 1m for the one corresponding to Fig. 4. For the remaining eight experiments, the robot ended within 3m of home position, which is less than 5% of the total distance traveled.

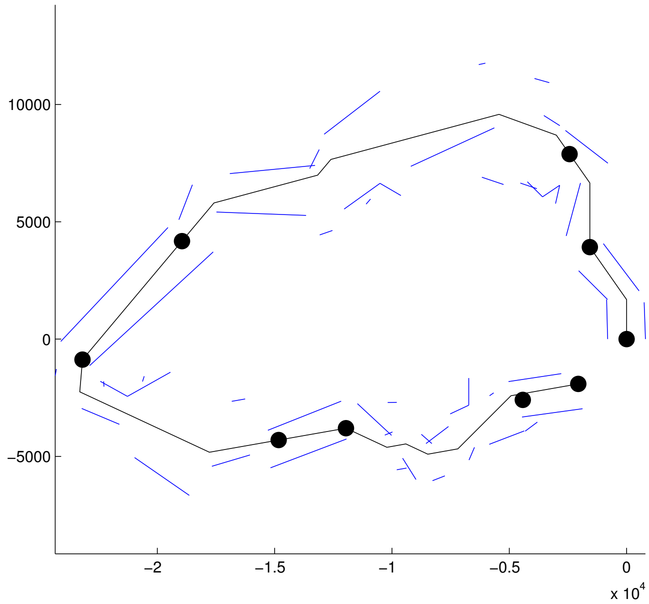

Fig. 4. ASRs computed in the homeward journey in the second experiment.

The robot's location in the physical environment during the homeward journey can be seen in Figures 5 and 6.

4.2 Orientation

In the homeward journey, the robot estimates where it is in the cognitive map computed during the outward journey. As such, it can use that position to estimate its orientation to home from its current position using the information contained in the map computed during the outward journey.

The robot can estimate the direction to the home at any intermediate position. Four randomly selected positions were chosen in each of the maps shown in Figures 3 and 4, and the estimated home direction was compared to the actual direction to the home position in the real world. The results are visualized in Fig. 5 (corresponding to Fig. 3) and Fig. 6 (corresponding to Fig. 4), which show a map of the real environment containing the path the robot actually took to return home. The estimated direction to home is depicted as a short arrow, the correct one as a long arrow. The estimated and correct angles with respect to the coordinate system of the map computed during the outward journey are given as well.

It can be observed that the direction estimate is fairly accurate; the accuracy usually decreases the longer the robot travels, which is due to error accumulation of the odometry measurements and drift.

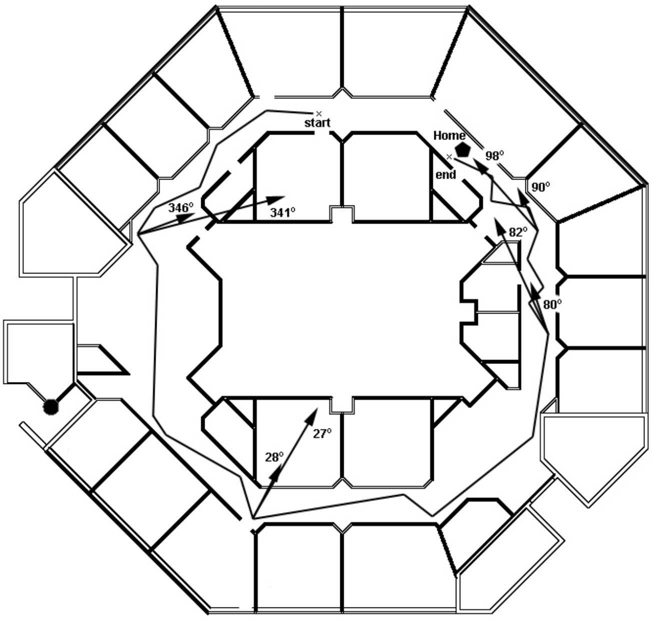

Fig. 5. Robot's estimations of home position (indicated by the shorter arrows) at four randomly selected positions during the first homeward journey. The longer arrow shows the correct orientation.

5 Discussion

The results we obtained for both experiments are surprisingly accurate. The robot did not go astray. It is interesting to note that in an earlier experiment [16], we used the following strategy:

1. Do not compute ASRs in the homeward journey.
2. Use the ASRs computed in the outward journey in a reverse order.
3. Measure the length of the ASR that the robot thinks it is in and travel similar distances to reach the end of that ASR.
4. Search for the entrance to the next ASR. If the next ASR is on its left, turn left. Otherwise turn right.

The robot performed less well using that strategy. This might argue well for the importance of computing local ASRs every time one enters a new local environment.

Our use of a robot with sonar sensors provides us with an opportunity to investigate the significant use of distance information in cognitive mapping. The surprisingly good initial results lead us to question whether there is ever a need to immediately combine different networks of ASRs. With hindsight, it now appears there are good reasons

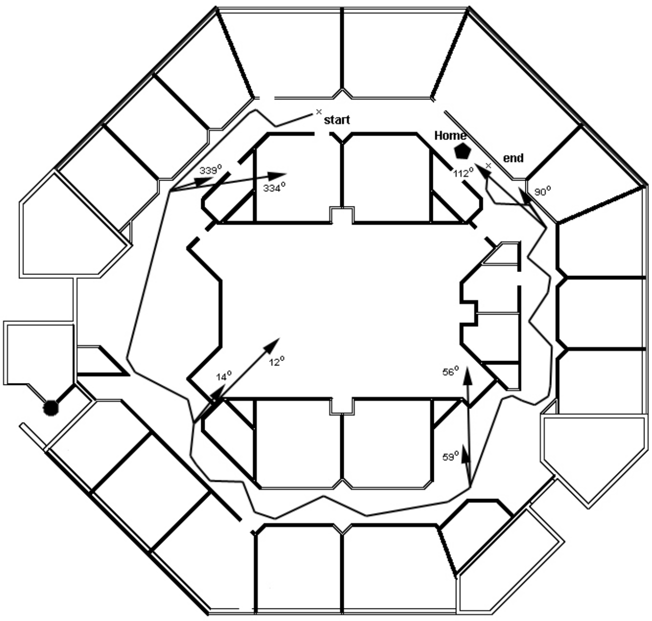

Fig. 6. Robot's estimations of home position during the second homeward journey.

that we do not want to. One reason is that if the environment is a simple environment, there is not much information to be gained from merging them. If the environment is complicated, then it is unclear an updated ASR is useful at all. Research into humans' ability to integrate paths such as that of Ishikawa and Montello [17] would help us in the future to refine our studies here.

The fact that the robot does not forget any of the ASRs along the way might help to explain the accurate orientation ability of the robot. Nonetheless the robot still needs to estimate its position fairly accurately in order to get a good orientation calculation. In the future, it would be interesting to explore how the robot might use orientation information to compute a short cut back home. It would also be useful to investigate means to orient if the network is not well-connected (i.e. with some ASRs missing, for example).

Finally, just like animals have shown to use ingenious algorithms to find its way, we could explore in the future many more different algorithms that exploit the use of a cognitive map to find its way.

Acknowledgments

This research project is supported by a grant from the Auckland University of Technology.

References

1. Yeap, W.K., Jefferies, M.E.: Computing a Representation of the Local Environment. Artificial Intelligence **107**(2) (1999) 265–301
2. Yeap, W.K.: Emperor AI, Where is your new mind? AI Magazine **18**(4) (1997) 137–144
3. Tolman, E.C.: Cognitive Maps in Rats and Men. Psychological Review **55**(4) (1948) 189–208
4. Golledge, R.G.: Human wayfinding and cognitive maps. In Golledge, R.G., ed.: Wayfinding Behavior: Cognitive Mapping and Other Spatial Processes. John Hopkins University Press, Baltimore (1999)
5. Chown, E., Kaplan, S., Kortenkamp, D.: Prototypes, location, and associative networks (plan): Towards a unified theory of cognitive mapping. Cognitive Science **19**(1) (1995) 1–51
6. Kuipers, B.: The spatial semantic hierarchy. Artificial Intelligence **119** (2000) 191–233
7. Poucet, B.: Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. Psychological Review **100** (1993) 163–182
8. Kortenkamp, D.: Cognitive Maps for Mobile Robots: A Representation for Mapping and Navigation. PhD thesis, University of Michigan (1993)
9. Jefferies, M.E., Cree, M., Mayo, M., Baker, J.T.: Using 2D and 3D landmarks to solve the correspondence problem in cognitive robot mapping. In: Proc. International Conference on Spatial Cognition IV 2004: Reasoning, Action, Interaction. Volume 3343 of Lecture Notes in Artificial Intelligence., Berlin, Springer-Verlag (2004) 434–454
10. Kuipers, B., Modayil, J., Beeson, P., MacMahon, M., Savelli, F.: Local Metrical and Global Topological Maps in the Hybrid Spatial Semantic Hierarchy. In: Int. Conf. on Robotics and Automation. (2004) 4845–4851
11. Ballard, D., Brown, C.: Computer Vision. Prentice-Hall Inc., Englewood Cliffs, New Jersey (1982)
12. Yeap, W.K.: Towards a Computational Theory of Cognitive Maps. Artificial Intelligence **34**(3) (1988) 297–360
13. Schmidt, J., Wong, C.K., Yeap, W.K.: A Split & Merge Approach to Metric-Topological Map-Building. In: International Conference on Pattern Recognition (ICPR). Volume 3., Hong Kong (2006) 1069–1072
14. Wong, C.K., Yeap, W.K., Schmidt, J.: Computing a Network of ASRs Using a Mobile Robot Equipped with Sonar Sensors. In: International Conference on Robotics, Automation, and Mechatronics (RAM), Bangkok, Thailand (2006) 57–62
15. Wong, C.K., Schmidt, J., Yeap, W.K.: Using a Mobile Robot for Cognitive Mapping. In: International Joint Conference on Artificial Intelligence (IJCAI), Hyderabad, India (2007) accepted for publication.
16. Wong, C.K., Yeap, W.K., Sapiyan, M.: A mobile robot that maps naively but plans intelligently. In: Proceedings of the Artificial Intelligence and Application Conference. (2005) 562–567
17. Ishikawa, T., Montello, D.: Spatial Knowledge Acquisition from Direct Experience in the Environment: Individual Differences in the Development of Metric Knowledge and the Integration of Separately Learned Places. Cognitive Psychology **52**(2) (2006) 93 – 129