

#

Using a Mobile Robot to Test a Theory of Cognitive Mapping

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

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

Abstract. This paper describes using a mobile robot, equipped with some sonar sensors and an odometer, to test navigation through the use of a cognitive map. The robot explores an office environment, computes a cognitive map, which is a network of ASRs [33, 34], and attempts to find its way home. Ten trials were conducted and the robot found its way home each time. From four random positions in two trials, the robot estimated the home position relative to its current position reasonably accurately. Our robot does not solve the simultaneous localization and mapping problem and the map computed is fuzzy and inaccurate with much of the details missing. In each homeward journey, it computes a new cognitive map of the same part of the environment, as seen from the perspective of the homeward journey. We show how the robot uses distance information from both maps to find its way home.

#.1 Introduction

A cognitive map is more than just a representation of the physical environment traversed. It should include, among other things, one's own experiences in the environment. For humans, this includes much of one's high-level interpretations of the environment itself. For this reason,

roboticists seldom refer to the map they compute as a cognitive map. Yet, a mobile robot's ability to move about autonomously and sense its environment by various means implies we should use a robot for testing different theories about cognitive mapping. However, researchers interested in developing computational theories of cognitive mapping rarely test their theories on a mobile robot, at least not initially, because implementing such a theory on a mobile robot is not a straightforward task. A few attempts have been made recently [18, 21, 22].

When discussing cognitive maps, one is concerned with epistemic and semantic issues. This is true even at the level close to perception. To illustrate, consider the notion of an egocentric representation of space. Roboticists would think of an egocentric representation of space as the equivalent of any representation computed from the viewpoint of a robot. Furthermore, its transformation to an object-centred representation is nothing more than a straightforward mathematical transformation. The egocentric representation of space becomes a special case of an object-centred representation, where the object is the viewer's body. Such a mathematical interpretation of an egocentric view ignores the fact that a view relates how a cognitive agent perceives its immediate space.

Philosophers such as Evans [10] and Campbell [2] have argued that the importance of an egocentric view is not that the representation is centred on the self but rather that the representation is action-guiding. Hence, although the representation is described mathematically as centring upon the self, the self is not identified in it. They argued that the actions afforded in each view are what matter most. Similarly, when humans perceive an object-centred representation of space, what is important is the realisation of the self as one of the objects in it and that space exists in an absolute sense. Space becomes something an individual moves into and space contains other objects. Interestingly, and as Pacherie [26] noted, since the self becomes one of the objects in such a representation, it could also then be used as its centre.

Out of a perceptual space, which is a relative view of space, animals (especially humans) are able to conceive space as absolute. Yeap [33] has argued that the cognitive mapping process should begin with the latter, which he referred to as an Absolute Space Representation (or ASR in short). However, investigations to date into the nature of ASR computations [16, 19, 20, 23] have focused primarily on its physical aspects i.e. information about an ASR that can be derived at a perceptual level. Such information includes its shape, its boundary surfaces, surfaces inside or outside the ASR, exits, etc. Such ASRs are then shown to be interconnected as a network of traversable regions, thus forming a cognitive map of the environment.

From a roboticist's point of view, computing a (cognitive) map in this way is but another method of partitioning the environment into traversable regions. The role such a map plays as a cognitive map is little realized. This is particularly true if the (cognitive) map is then used like a cartographic map; successful use depends very much on the metric accuracy of the ASRs and less so on any kind of heuristic reasoning that animals (especially humans) aptly apply. So, where lies the cognitive sense of the map computed?

Just like the mathematical notion of egocentric representation, computing ASRs is indeed equivalent, mathematically, to the partitioning of space into separately identifiable regions. Since the introduction of the first mobile robots such as Shakey [25], roboticists have devised many such algorithms (see Section 1 of this book on robot mapping and for some classic examples see [3, 8, 28]). However, just like an egocentric view of space, what matters most in an ASR's computation is not the partitioning of space per se into separately identifiable regions for the individual to traverse between them, but the formation of a foundation for developing a much higher level description of space. If returning to a given place were important, animals could, and have been shown to, evolve more direct algorithms for arriving home (such as the various methods for path integrations, see [7, 9]). An ASR affords the development of one of the most basic elements in cognitive maps, namely the notion of a place.

The word "absolute" in the acronym, ASR, emphasizes the existence of the space itself into which one has entered. One computes a new ASR as soon as one enters a new local space, presumably to quickly identify the new local space into which one has just entered. Depending on the cognitive capabilities of the individual species, the new ASR would rapidly be transformed from a space-sense to a place-sense i.e. from a purely perceptual sense to an increasingly conceptual one. For example, in humans, once the exits and boundaries are identified, the cognitive mapping process does not stop there. Boundaries could be interpreted as closed doors (which imply they are actually exits) and gaps that are perceived as exits might be interpreted as gaps between, say, two pillars inside the ASR (which imply they are not exits). Surfaces will be interpreted as objects and their functional significance realised. Objects could become landmarks. Events might be unfolding which then draw the attention of the individual further away from attending to the physical qualities of its environment. The notion of a place which begins with a single ASR could become a collection of ASRs, the map a network of places.

Ideally, we have to show how ASRs are turned into places and the map used to solve various spatial tasks. Computing place representations

provides a far more useful basis for reasoning about one's environment than computing the physical shape of ASRs alone and can only come with powerful reasoning capabilities, which, for now, the robot lacks. A blind rat can find its way in a maze. What we could investigate now are the different (cognitive mapping) ways in which the map could be used by the robot. Just like birds have a mapping process different from a rat, and rats from humans, robots can have their own cognitive mapping process¹. By studying the cognitive mapping process of the different species (robots included), we gain better insights into the nature of the cognitive mapping process.

What would a cognitive mapping process of a robot be, given the sensors it has? Note that existing works (e.g. see Jefferies et. al. in Chapter X) have shown how a robot can compute a network of ASRs and use it to navigate in its environment. The ASRs computed are reasonably precise in metric terms and none is missing from the robot's memory. Such a robot is analogous to a cognitive agent navigating in a familiar part of the environment (where things are remembered fairly precisely).

In this paper, we show another example of how a robot computes and uses cognitive maps. In particular, we ask what would the process be like if, at the end of an initial exploration of a new environment, the robot does not have a well-formed network of ASRs? What if some of, or all the ASRs computed are incomplete and contain inaccurate information? Presumably, the robot could not localise itself in the environment using the map computed. What should it do to find its way home? Note that the ASRs computed when finding its way home could be very different from those computed during the initial exploration. How are the different maps, one in the memory and the other currently experienced, being used? How do we set up an experiment with our robot to investigate some of these questions?

The remainder of this paper describes an experiment with a mobile robot doing cognitive mapping. Our mobile robot is equipped with some sonar sensors and an odometer. Sonar sensors provide very inaccurate and unreliable measurements. They were deliberately chosen so that the network of ASRs computed cannot function as an accurate map. Instead, these ASRs serve more like fuzzy memories of places visited. In the experiment, the fuzziness of the ASR is, of course, due to the poor sensors used. It is meant to simulate the typical kind of fuzzy memory of a place humans recall about the new environment they have just visited.

¹Strictly speaking, a robot, not being a cognitive agent, cannot have a cognitive mapping process. However, rather than being verbose and say "simulating a cognitive mapping process", we will simply say the robot has a cognitive mapping process.

Section #.2 describes how we set up our robot to explore its environment and compute fuzzy ASRs. Section #.3 describes the problem faced by our robot. Section #.4 describes the algorithms the robot used and the results of our experiment. Section #.5 concludes with a discussion of cognitive and robotic mapping in the context of our experiments.

#.2 The Robot and Its Cognitive Map

The robot we used is a Pioneer 2 robot from *ActivMedia* and it came with a ring of 8 sonar sensors. The robot was positioned somewhere in the corridor in an office environment and was allowed to explore the environment until it was told to stop. No modifications to the environment were done. Everything already in the environment (such as rubbish bins, flower pots, cabinets, etc.) remained where they were and doors leading into offices were left as they were found, closed or open, at the time of the experiment.

The environment used and one of the paths the robot took is as shown in Figure #.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 eight sensors but information about the environment is sensed only via the two side sensors. That is, the robot uses its eight sensors to decide where to move next but only its two side sensors to gather information about the environment. The exploration algorithm used is 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 but maintain a forward-going direction

The details of our new algorithm for computing ASRs for the experiments conducted here can be found in [29, 32]. Briefly, the key ideas underlying our new algorithm are:

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);




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

2. The important exits found in a path are the exits at both ends of it (i.e. given the poor sensing, it 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 – 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 reasonably sized boundary is obtained.
4. An ASR computed for a path represents an ASR computed from a single view of the robot. The next step is to merge or split ASRs obtained from individual paths into the final network of ASRs for the environment experienced.

Figure #.2 shows the final ASRs computed for the journey shown in Figure #.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. Note that in the area marked A the robot is shown moving through that part of the environment in a single path. That single ASR was split into three ASRs for that part of the environment. In the area marked B, it shows the robot moving through it using five paths. These five ASRs are later merged to form the final single ASR. Such merging and splitting is done in step 4 of the algorithm as discussed above.




Fig. #.2. ASRs computed for the journey as shown in Figure #.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.

#.3 The Experiment




Fig. #.3. ASRs computed in the homeward journey. (0,0) indicates the starting position of the robot.

You could imagine that our robot is a blind “rat” with a special sense. It could stretch out a pair of imaginary arms from its side to infinity, or until it “believes” its arm touches an object. Consequently, what it senses might not be correct. The object might or might not be there. In this manner, it gathers information about the shape of the environment as it moves down a path. ASRs are then computed from the information gathered. Without vision, ASRs are computed after the robot has left that part of the environment.

Like a typical rat experiment, our experiment is to let our robot wander freely between two points in an office environment, imagining them to be the home and food locations in a maze. The robot must then find its way home and the question is how. The robot has a cognitive map of its environment computed during its outward journey. During its homeward journey, it will compute another cognitive map. Figure #.3 shows an example of a network of ASRs generated during one of its homeward journeys. The network is not the same as that computed during the outward journey. In particular, nine ASRs are computed instead of ten and the sixth ASR computed in the outward journey is now perceived as three ASRs in the homeward journey.

Our robot cannot confidently re-compute incoming ASRs and match them with those in memory. Thus, it does not solve the widely accepted problem among robotics researchers, namely the simultaneous localization and mapping problem (famously known as the SLAM problem – see for example, [6]). Our robot is at a very early stage of (cognitive) mapping of its environment. In many studies of cognitive mapping (for example, see [13]), much has been said about the use of distance and direction information in the process. We implemented an algorithm for our robot to find its way home that uses distance information implicit in each cognitive map. Given the cognitive maps computed, how good is the robot's sense of direction? We also implemented an algorithm for the robot to estimate its orientation in its current position to the home position.

#.4 Implementations and Results

Section #.4.1 presents the implementation of our home-going algorithm that makes use of the ASR distance traveled as opposed to the actual distance traveled by the robot in its zigzag moves home. A total of ten experiments were conducted using this algorithm; the results for two of them are described in more detail. Section #.4.2 presents an algorithm for the robot to estimate its orientation towards home and the result is compared with its actual orientation from where it is physically located in the environment.

#.4.1 Going Home

The algorithm for returning home is:

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 for how far the robot has re-traced its steps towards home. The robot stops when it believes it has completely retraced its steps.

Figure #.2 shows the cognitive map computed by the robot for its outward journey. Ten experiments were conducted using our “Going Home” algorithm for the robot to find its way home. Figures #.3 and #.4 show two cognitive maps computed in two different attempts to go home.




Fig. #.4. ASRs computed in the homeward journey in the second experiment. (0,0) indicates the starting position of the robot.

We measured the distance between the robot's final position and the real home position. For the two experiments presented in the figures, the robot was 1.5m short of the true home position in the experiment corresponding to Figure #.3, and 1m short for the experiment corresponding to Figure #.4. For the remaining eight experiments, the robot ended within 3m of home, which is less than 5% of the total distance traveled. The robot's positions in the physical environment during the homeward journeys shown in Figures #.3 and #.4, can be seen in Figures #.5 and #.6 in the next section.

#.4.2 Orientation




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.

During the homeward journey, the robot estimates where it is in the cognitive map it computed for the outward journey. It estimates its orientation to home from its current position using the information contained in the outward journey's map. This part of the experiment is to investigate how accurate is the robot's sense of home's direction.

The robot can estimate the direction to 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 [from each position](#) was compared to the [real world direction](#).




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

The results are visualized in Figures #.5 (corresponding to Figure #.3) and #.6 (corresponding to Figure #.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 estimated is fairly accurate; it has not been affected by errors due to odometer measurements and drift.

#.5 Discussion

Our robot has a very limited capability for sensing its environment. It uses cheap sonar sensors (as opposed to the more advanced sonar sensors used in [1, 5, 27] or those with powerful pre-processing software [30]). Nonetheless, even with such limited sensing, we have shown that it is possible to implement an algorithm for our robot to compute a network of ASRs. The notion of an ASR is versatile and is not restricted to having accurate or powerful sensors. Our work thus highlighted one significant difference between partitioning the environment into traversable regions (robot mapping) and into ASRs (cognitive mapping). Robot mapping is more concerned with dividing a large-scale space into smaller ones and their physical qualities whereas cognitive mapping is about the identity and character of each local space visited. For cognitive mapping, the more versatile notion of space is preferred.

We have implemented a very basic algorithm for the robot to find its way home, namely exploit ASR-distance traveled to re-trace its movements to return home. Much has been discussed with respect to the use of distance information in cognitive mapping. For example, numerous experiments with chickens and pigeons have shown that they are able to use both absolute and relative distance in their search for food (see [14] for an example of recent work). Experiments with bees and ants have shown that they can perform internal calculations of the distance and direction traveled to perform path integration (see [7, 9] for a general discussion). Most of these experiments were concerned with the actual distance traveled and how the individual species deal with the errors in their measurements, as do most work on robot mapping to date. Using our robot, we have shown another way of using distance information, namely ASR-distance traveled as opposed to actual distance traveled. The method appeared to be useful.

ASR-distance is obtained from the shape of the ASR computed. In the past, there has been scant evidence that humans/animals do pay attention to the shape of each local environment (or, in our terminology, ASR) very early on in their initial exploration of a new environment [4, 11, 12, 15, 17, 24]. However, the debate has now intensified and this is especially true in the animal literature where the problem is commonly referred to as geometry in animal spatial behavior (see Cheng in Chapter X). In a relocation task using a box-shaped environment, the principal axes of the environment appear to be most useful. However, Cheng questioned the general applicability of the principal axes and suggested other ways of utilizing the shape of “ASRs” computed. Our work here emphasized yet another possibility, namely using a straight line distance between exits of interests in an ASR.

Two remarks are worth making regarding the surprisingly good results obtained in our experiment. First, although our robot was allowed to wander on its own during all the trials, it managed not to enter any of the rooms. Consequently, the robot appears to be constantly moving forward along the corridor and this might have accounted for much of the success of the experiment. This was not planned. It would be interesting to see how the resulting ASRs would be if the robot enters the middle room and follows a path such as that shown in Figure #.7.

The ASR algorithm would have to be made more powerful so that it could reason about the overall shape of the ASRs computed. If no other kinds of sensors are used, this robot would not be able to learn much about its environment; it could not identify any objects in it. In the future we are planning to add other kinds of sensors (e.g. compass) to our robot to investigate how the extra information made available will enhance the robot’s reasoning about its environment.

Second, it is interesting to note that in an earlier experiment [31], the following strategy is used:

1. Do not compute ASRs during the homeward journey.
2. Use the ASRs computed for the outward journey in 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. Once it believes it has reached the end of the ASR, search for the entrance to the next ASR. If the next ASR is on its left, turn left. Otherwise turn right.




Fig. #.7. Robot's possible navigation inside a room– Will it compute a single ASR for the room or multiple ASRs?

The robot did not perform as well using the above strategy (4 out of 6 trials were successful). However, it does cause the robot to exhibit some interesting behavior at each decision point (step 4 in the above algorithm). In trying to find the exit, the robot makes small turns and movements, appearing to be cautious in its search for an exit. Figure #.8 shows an example of a path taken by the robot using the above algorithm.

Given the current restricted paths through the environment and the small number of trials conducted, it is not interesting comparing the performance of the two algorithms. Rather, what is interesting is to observe that the two different algorithms represent two different approaches to using and updating a cognitive map. The first approach is to always compute an ASR from the input and then extract information from it for comparisons or updating with those held in one's cognitive map. The other is to directly use information from the input with those held in one's cognitive map.




Fig. #.8. The path the robot took on its way home. The points marked X, Y, and Z are critical decision points in the journey home.

We need to investigate many more different strategies before we can understand how the different strategies interact in a cognitive mapping process.

Finally, we implemented an algorithm for the robot to tell us where it believes its goal is from its current position in its cognitive map. We compare that with the actual orientation of the goal from the robot's physical location in the building. The fact that the robot does not forget any of the ASRs along the way might help to explain the robot's ability to accurately orient itself. In the future, we would like to explore how the robot might use the orientation information to compute a short cut to home. It would also be useful to investigate means to orient itself if the network is not well-connected (i.e. with some ASRs missing, for example).

#.6 Acknowledgments

This research project is supported by a grant from the Auckland University of Technology. We are grateful to Ewing Caldwell for his comments on an earlier draft.

References

1. Barshan B, Kuc R (1990) Differentiating Sonar Reflections from Corners and Planes by Employing an Intelligent Sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12:6:560-569
2. Campbell J (1993) The Role of Physical Objects in Spatial Thinking. In: Eilan N, McCarthy R, Brewer B (Eds) *Spatial Representation: Problems in Philosophy and Psychology*, Oxford: Blackwell Press, pp 65-95
3. Chatila R (1982) Path Planning and Environment Learning in a MobileRobot System. In: the Fifth European Conference on Artificial Intelligence. Paris, France, pp 211 – 215
4. Cheng K (1986) A Purely Geometric Module in the Rat's Spatial Representation. *Cognition* 23: 149-178
5. Chong KS, Kleeman L (1999) Mobile Robot Map Building for an Advanced Sonar Array and Accurate Odometry. *Int J Robotics Research* 18:1:20-36
6. Choset H, Nagatani K (2001) Topological Simultaneous Localization and Mapping (SLAM): Toward Exact Localization without Explicit Localization. *IEEE Transactions on Robotics and Automations*, 17:2:125-137
7. Cornell EH, Heth CD (2004). Memories of travel: Dead reckoning within the cognitive map. In: GL Allen (Ed.) *Human Spatial Memory: Remembering Where*. Mahwah, NJ: Lawrence Erlbaum Associates, pp 191-215
8. Crowley JL (1985) Navigation of an Intelligent Mobile Robot. *IEEE J Robotics Research* 1:1:31-41
9. Etienne AS, Maurer R, Georgakopoulos J, Griffin A (1999) Dead Reckoning (Path Integration), Landmarks, and Representation of Space in a Comparative Perspective. In: Golledge RG (Ed.) *Wayfinding Behavior: Cognitive Mapping and other Spatial Processes*. Baltimore: John Hopkins University Press, pp 197-228
10. Evans G (1982) *The Varieties of Reference*. Oxford: Oxford University Press

11. Gallistel CR (1989) Animal Cognition: The Representation of Space, Time and Number. *Annual Review of Psychology* 40:155-189
12. Gallistel CR, Cramer, AE (1996) Computations on Metric Maps in Mammals: Getting Oriented and Choosing a Multi-Destination Route. *J Experimental Biology* 199:211-217
13. Golledge RG (1999) Wayfinding Behavior: Cognitive Mapping and other Spatial Processes. Baltimore: John Hopkins University Press
14. Gray ER, Spetch ML, Kelly DM, Nguyen A (2004) Searching in the Center: Pigeons Encode Relative Distance from Walls of an Enclosure. *J Comparative Psychology* 118:1:113-117
15. Hermer L (1997) Internally Coherent Spatial Memories in a Mammal. *Learning and Memory* 8:1743-1747
16. Hill Jr. RW, Han C, Van Lent M (2002) Applying Perceptually Driven Cognitive Mapping to Virtual Urban Environments. In: National Conference on Artificial Intelligence, AAAI press, Edmonton, Canada, pp 886-893
17. Huttenlocher J, Newcombe N, Sandberg EH (1994) The Coding of Spatial Location in Young Children. *Cognitive Psychology* 27:115-147
18. Jefferies ME, Cree M, Mayo M, and Baker JT (2005) Using 2D and 3D Landmarks to Solve the Correspondence Problem in Cognitive Robot Mapping. In: Freksa, C, Knauff M, Krieg-Brückner B, Nebel B, and Barkowsky T (eds) *Spatial Cognition IV. Reasoning, Action, and Interaction: International Conference Spatial Cognition*. Frauenchiemsee, Germany, pp 434-454
19. Jefferies, ME, Yeap WK (1998) Representing the local space qualitatively in a cognitive map. *Proceedings of the Cognitive Science Conference*. Madison, USA, 525-530
20. Jefferies ME, Yeap WK (2001) The utility of global representations in a cognitive map. In: *Conference on Spatial Information Theory*. Morro Bay, USA, pp 233-246
21. Kortenkamp, D (1993) *Cognitive Maps for Mobile Robots: A Representation for Mapping and Navigation*. PhD Thesis, University of Michigan
22. Kuipers B, Modayil J, Beeson P, MacMahon M, Savelli F (2004) Local Metrical and Global Topological Maps in the Hybrid Spatial Semantic Hierarchy. In: *IEEE International Conference on Robotics and Automations*, New Orleans, USA, pp 4845-4851
23. Lavieri KR, Peterson GL (2004) Cognitive Robot Mapping with Polylines and an Absolute Space Representation. In: *IEEE International Conference on Robotics and Automations*, New Orleans, USA, pp 3771-3776

24. Margules J, Gallistel CR (1988) Heading in the Rat: Determination by Environmental Shape. *Animal Learning and Behaviour* 16:4:404-410
25. Nilsson NJ (1969) A Mobile Automaton: An Application of Artificial Intelligence Techniques. *Proceedings of the First International Joint Conference on Artificial Intelligence*, Washington, pp 509-519
26. Pacherie E (2002) Perspectival and Objective Representations of Space and Time. *Proceedings of the 4th European Congress of Analytic Philosophy*, Lund, Sweden
27. Peremans H, Audenaert K, Van CJM (1993) A High-resolution Sensor Based on Tri-aural Perception. *IEEE Trans. Robot Automation*, 9:1:36-48
28. Rueb KD, Wong AKC (1987) Structuring Free Space as a Hypergraph for Roving Robot Path Planning and Navigation. *IEEE Transactions of Pattern Analysis and Machine Intelligence* 9:2:263-273
29. Schmidt J, Wong CK, Yeap WK (2006) A Split and Merge Approach to Metric-Topological Map-Building. Accepted to the 18th International Conference on Pattern Recognition, Hong Kong
30. Tardos JD, Neira J, Newman PM, Leonard JJ (2002) Robust Mapping and Localisation in Indoor Environments using Sonar Data. *Int J Robotics Research*, 21:4:311-330
31. Wong CK, Yeap WK, Sapiyan M (2005) A Mobile Robot that Maps Naively but Plans Intelligently. In: Hamza MH (ed) *Artificial Intelligence and Application*, ACTA Press, Innsbruck, pp 562-567
32. Wong CK, Yeap WK, Schmidt J (2006) Computing a Network of ASRs Using a Mobile Robot Equipped with Sonar Sensors. Accepted to the IEEE International Conference on Robotics, Automation and Mechatronics, Bangkok
33. Yeap WK (1988) Towards a Computational Theory of Cognitive Maps. *Artificial Intelligence* 34:3:297-360
34. Yeap WK, Jefferies ME (1999) Computing a Representation of the Local Environment. *Artificial Intelligence* 107:2:265-301