COMPUTING A QUALITATIVE REPRESENTATION FOR THE LOCAL SPACE

MARGARET E JEFFERIES and WAI-KIANG YEAP
Artificial Intelligence Laboratory
University of Otago, Dunedin, New Zealand
e-mail: jefferies_me@cs.otago.ac.nz

ABSTRACT

Yeap and Jefferies [1] show how a description of each local space is computed and by realising how these local spaces are connected, a cognitive map of one's environment will emerge. However, roboticists have shown that it is difficult to compute such a map with precise metric information. In this paper, we show some preliminary ideas as to how a qualitative map could be constructed out of local spaces. The trick is to devolve each representation to make it less precise and then slowly learn more about each space so that it later becomes unique.

INTRODUCTION

Much of the earlier work on qualitative representations of space are concerned with providing a suitable calculus for representing and reasoning with spatial entities. Examples of such work include: [2-6]. In this study, we are concerned with how a qualitative representation of space is derived from what is perceived and how such a representation can later become a useful cognitive map.

Humans' immediate perception of the environment is precise in the sense that they could perform a variety of spatial manipulation tasks. Yet, rarely are they able to reproduce from memory an exact description of the places they have visited. We do not often remember exactly where objects are when we can't physically see them and yet we are able to make good use of the vague and imprecise memories we have for our environment. For whatever reasons, it is thus clear we transform from a precise description to an imprecise one. When doing so, what effect does this have on the cognitive mapping process?

We have been studying the problem of computing a cognitive map using visual information derived from a simulated 2D environment [1, 7-10]. Recent AI models of cognitive mapping include the work of Engelson [11] and Kortenkamp [12]. More recently, we presented an algorithm which successfully computed a description of each local space (which we called an absolute space representation, or ASR) as a simulated robot moves through a predefined path [1]. By devolving each ASR so that it becomes an approximate description of the local space, we get a qualitative map. In this paper, we describe some initial experiments with such a map.

QUALITATIVE LOCAL SPACE REPRESENTATIONS

Determining the ongoing nature of spatial memories when they are no longer receiving immediate feedback from the environment is not easy. Studies which examine this problem are mostly concerned with the manner in which the representations are distorted and their significance altered once they are merged into the wider "picture in the head". Variables such as size, distance and location are often systematically distorted by containment relations and the significance of an object as compared with others it is related to [13-16]. But these modifications result from some top-down processing, i.e. the input to the process isn't only what has been computed bottom-up from the senses but also includes the results of earlier computations, often higher-level representations which are conceptually more sophisticated.

Our concern at this stage is only with what can be computed bottom-up from the senses. It is our contention that the initial representation computed for a local space is computed for the viewer's immediate needs, to provide a locus for the objects surrounding the viewer, and the activities which involve these objects. But while much of the detail is forgotten or goes unnoticed one can still remain cognizant of the local space for a long time after it was occupied. To encapsulate a local space in this way would only require representing its extent in very rough terms, but the resulting representation could still provide a useful framework for reasoning about the local space and could then become more elaborate as the viewer's familiarity with the environment increases.

We compute such a representation by devolving the initial representation computed into a rectangle which roughly approximates its extent. A straightforward algorithm is used - points on the surfaces forming the boundary of the ASR are sampled to firstly find a good length for the rectangle and then the length itself is sampled to find a good width. We call this representation a fuzzy ASR. Figure 1 (a) shows an initial ASR computed, its surfaces are labelled s1 – s5 and its exits e1 – e4 (for a detailed description of this algorithm, see [1]. The fuzzy ASR computed from this ASR is shown in Figure 1 (b). No claims are made as to the cognitive plausibility of our method. In reality many processes would be operating to modify the origi-

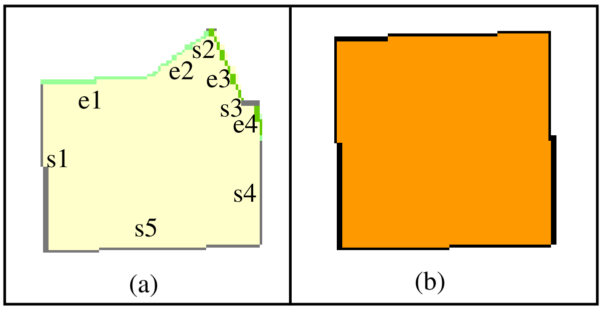

Figure 1. (a) An exact ASR computed while the viewer occupied the local space. (b) the fuzzy ASR description which devolves from the ASR in (a)

nal ASR and we cannot claim to fully understand these. This is but one method for producing a fuzzy ASR. There will be many, many more.

The real significance of the fuzzy ASR for our computational theory is the manner in which the representation is able to be used to structure the cognitive map, however poorly. The fuzzy ASR does not comprise actual surfaces or exits, it merely represents a portion of space once occupied by the viewer. But one would expect the viewer to remember some of the connections to neighboring spaces, confused though they may be. Thus we retain the connections to neighboring "ASRs", but only in the loosest sense. We conducted three experiments with the program by varying the amount of knowledge the viewer retained for the connections between ASRs. We thus showed how a fuzzy cognitive map might be structured and how useful such a map might be.

In the first experiment the viewer remembers how many exits there are in an ASR but no locational information is retained for them. For the fuzzy ASR in Figure 1 (b), for example, the viewer remembers just that there are four exits, e1, e2, e3 and e4. When the ASR is exited a connection is made to the ASR just entered but the viewer does not remember which exit was used. Our viewer has a very poor memory indeed! The outcome of this is a scenario often faced by humans - "I know I've been here before so which doorway did I use to get to..." Thus the information made explicit in a fuzzy-ASR comprises the rough extent of the ASR, the number of exits in the ASR and which neighboring ASRs have been experienced as connected to this one. The results of the experiment are displayed in Figure 2. Figure 2 (a) shows the portion of the environment traversed and Figure 2 (b) a cognitive map constructed from the "exact" ASRs computed for each local space visited. Note that although for display purposes the ASRs are laid out as if there is one global coordinate system, in reality this is not the case. Each ASR is independent of all others with its own local coordinate system, and the only links to other ASRs are through the exits used to traverse them. Thus the viewer knows exactly where each surface and exit in the ASR is located, and exactly which exits are used to connect to neighboring ASRs. The actual structure of the cognitive map is mostly route-like, except where previously visited ASRs are able to be recognized (see [9]). These parts of the map exhibit a more integrated structure. Figures 2 (c) and (d) convey the underlying structure of the fuzzy cognitive map more realistically but this is only practical for a small number of ASRs. In all the figures the ASRs are numbered in the order in which they are visited.

The fuzzy cognitive map constructed for the path in Figure 2 (a) would comprise:

fuzzy-ASR1 with four exits, connected to ASR 2
fuzzy-ASR2 with two exits, connected to ASR 3, ASR 1
fuzzy-ASR3 with five exits, connected to ASR 4, ASR 2
fuzzy-ASR4 with three exits, connected to ASR 5, ASR 3

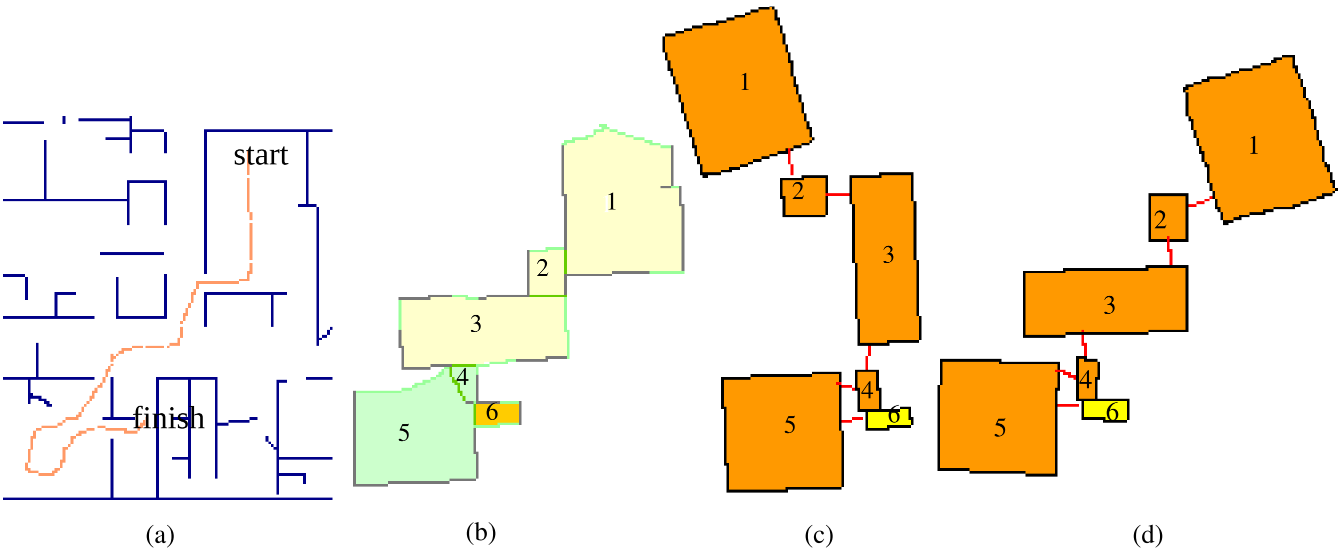

Figure 2. (a) The environment traversed, (b) a cognitive map computed from exact ASRs (c) the viewer's interpretation of a cognitive map computed from fuzzy ASRs where the viewer has no locational information for the exits. (d)the viewer's interpretation of a cognitive map constructed from fuzzy ASRs where the viewer knows on which side of the fuzzy ASR the exits are located but not their exact position.

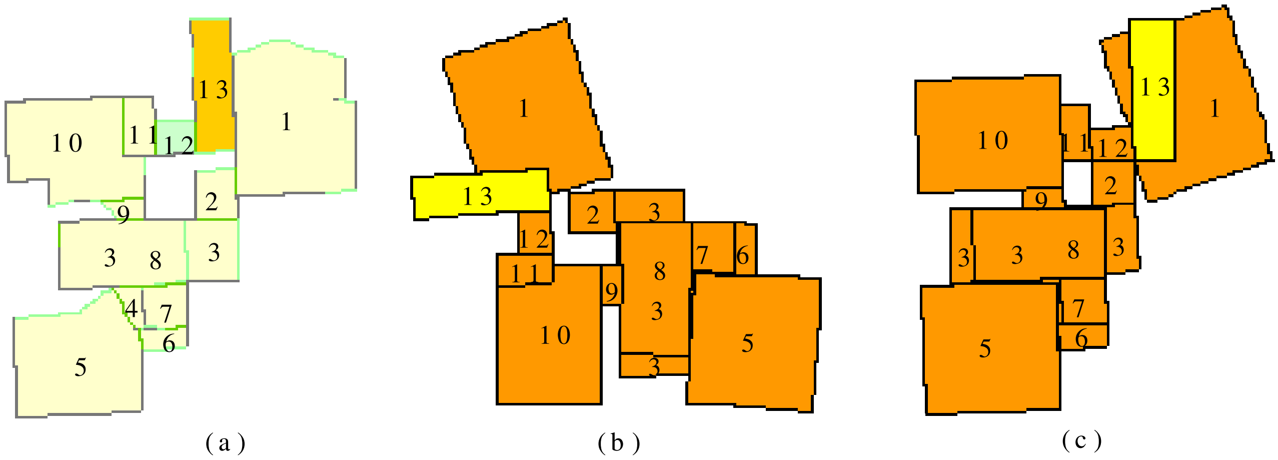

Figure 3.(a) a cognitive map computed from exact ASRs. (b) the viewer's interpretation of a cognitive map computed from fuzzy ASRs where the viewer has no locational information for the exits. (c) the viewer's interpretation of a cognitive map constructed from fuzzy ASRs where the viewer knows on which side of the fuzzy ASR the exits are located but not their exact position.

fuzzy-ASR5 with four exits, connected to ASR 6, ASR 4
fuzzy-ASR6 with three exits, connected to ASR 5

To demonstrate the usefulness of such a map, the viewer is told to repeat the journey from start to finish in its head. Figure 2 (c) demonstrates how confused a viewer making use of such a map could become. As the viewer imagines re-entering ASR 1, its knows from its fuzzy map that one of these exits leads into ASR 2 but not which one. The viewer randomly chooses an exit. The line emanating from the bottom of fuzzy ASR1, rather than its side, demonstrates that the viewer made an erroneous decision. It can be seen from the output from our computer simulations displayed in this figure that the errors made here result in rotation errors in the cognitive map and while they are not shown in this figure, translation errors are possible also.

In the second experiment we allowed the viewer to remember on which side of the fuzzy ASR the exits were located and thus on which side of a fuzzy ASR the connection to a particular ASR is located. Figure 2 (d) shows a viewer's attempt at using a fuzzy cognitive map constructed using this strategy. In ASR 1 the viewer recalls that ASR 1 connects to ASR 2 via an exit on the left side of ASR 1 and since there is only one such exit the correct choice is made. However, on the side of ASR 3 which connects to ASR 4 there are two exits. One leads directly into ASR 4 (see Figure 2 (b)) and one leads into an as yet unexplored region of the environment – this exit can be seen as the lighter shaded gap in the boundary directly adjacent to the exit into ASR 4 in Figure 2 (b). To visit ASR 4 from ASR 3 the viewer must choose between these exits and does so correctly (this time). If the incorrect exit had been chosen a translation error would have occurred. This is the case in Figure 3 (d) when the viewer does make a wrong decision on which exit leads from ASR 3 into ASR 4. See the paragraph which follows for a more detailed explanation.

Figure 3 shows the results of applying the strategies of both experiments to a longer traversal of the environment. Figure 3 (a) shows the cognitive map constructed from exact ASRs, (b) the cognitive map constructed when the viewer has no locational information for the exits in a fuzzy ASR, and (c) the cognitive map constructed when the viewer remembers which side of the fuzzy ASR the exits are on. Note that in all the situations, the viewer fails to recognize ASR 3 when the local space is re-entered from ASR 4 and a new ASR, ASR 8 is constructed. This is overlaid on top of ASR 3 only for display convenience. There is no such integration in the viewer's "head" and a one dimensional route-like structure is a better approximation of the actual structure of this part of the cognitive map. Figure 3(b) has the expected rotation error. A translation error occurs at about fuzzy ASR 3 in both Figure 3 (b) and Figure 3 (c). This is most noticeable in the way in which fuzzy ASR 8 in particular, has shifted in relation to fuzzy ASR 3 in the display. In deciding which exit leads from fuzzy ASR 3 into fuzzy ASR 4 the viewer selects an erroneous one which is on the same side of the fuzzy ASR as the correct one. In Figure 3 (c) it is just possible to make out a corner of fuzzy ASR 4 underneath fuzzy ASR 5. Unfortunately in Figure 3 (b) fuzzy ASR 4 is completely hidden.

The third experiment shows the viewer exploiting the structure of its fuzzy cognitive map to find its way back along a previously traversed path and failing to do so. The viewer firstly walked a path through its environment computing the fuzzy cognitive map displayed in Figure 4 (b), starting from the "home" ASR and ending at the ASR labelled "walk's end". As in Figure 3 (b), the display portrays how the viewer imagines the individual ASRs are connected when exits are chosen randomly. Figure 4 (a) shows the accurate map that would be computed for the same walk. The viewer remembered only that 10 ASRs had been visited in turn and therefore attempted to go back

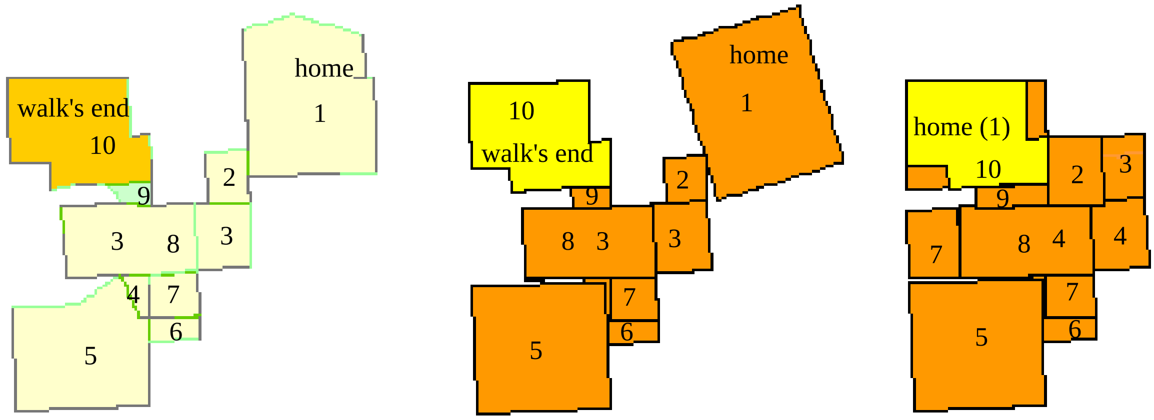

(a) (b) (c)

Figure 4 (a) a cognitive map computed from exact ASRs (b) the viewer's interpretation of a cognitive map computed from fuzzy ASRs where the viewer has no locational information for the exits. (c) the sequence of ASRs visited as the viewer tries to make its way home exploiting the structure of (b), i.e. tries to visit the ASRs in the reverse sequence to that in which they were originally computed.

through 10 ASRs. In Figure 4 (c) the viewer, on reaching ASR10, turns around and heads for home. With no contextual information available as a guide, and knowing only that one of ASR 10's exits leads into the next ASR along the path, i.e. ASR 9, the viewer randomly chooses an exit. It chooses correctly and enters ASR 9. Note that fuzzy ASR 9 in Figure 4 (b) is a different shape to the one computed in Figure 4 (c) because the ASR is entered from a different side. When the viewer is in ASR 8 the wrong exit is chosen to enter ASR 7, but this time realises it has reached a dead-end, returns to ASR 8 and then makes the correct choice. On entering ASR 5, the viewer is able to reach ASR 3 of Figure 4 (b) directly so when ASR 5 is exited the viewer thinks this ASR is ASR 4. ASR 2 of Figure 4 (b) is then entered as ASR 3. On exiting ASR 3 the viewer takes the wrong exit, choosing the one which will take it away from the "true" home ASR. Lastly the viewer ends up back where it started in the ASR at the end of the walk, but thinking the "home" ASR has been reached.

CONCLUSION

By carefully controlling the amount of information made available in a fuzzy ASR, it is hoped that future work will show how a cognitive map is learned.

REFERENCES

1. W.K. Yeap and M.E. Jefferies, *Computing a representation of the local environment*, AIM-35-97-1, (Computer Science Dept, University of Otago, 1997).
2. A.G. Cohn and N.M. Gotts, Representing spatial vagueness: A mereological approach, in *5th International Conference on Principles of Knowledge Representation and Reasoning (KR96)*, 1996, 230-241.
3. A.G. Cohn, D.A. Randell and Z. Cui, Taxonomies of logically defined qualitative spatial relations, *International Journal of Human-Computer Studies* 43(5-6), 1995, 831-846.
4. D. Hernandez, Reasoning with qualitative representations: exploiting the structure of space, in *The III Imacs International Workshop on Qualitative Reasoning and Decision Technologies - QUARDET '93*, 1993, 493-502.
5. D. Hernandez, E. Clementini and P. DiFelice, Qualitative distances, in *Spatial Information Theory- A Theoretical Basis for GIS*, Frank, A.U. and Kuhn, W., eds. (Springer-Verlag, 1995) 45-59.
6. B. Kuipers, A hierarchy of qualitative representation for space, in *10th International Workshop on Qualitative Reasoning about Physical Systems (QR-96)*, 1996.
7. W.K. Yeap, Towards a computational theory of cognitive maps, *Artificial Intelligence* 34, 1988, 297-360.
8. W.K. Yeap, P.S. Naylor and M. Jefferies, Computing a representation of the physical environment, in *First Pacific Rim International Conference on Artificial Intelligence*, 1990, 847-852.
9. W.K. Yeap, M. Jefferies and P. Naylor, An MFIS for computing a Raw Cognitive Map, in *Proceedings of the 12th IJCAI*, 1991, 373-378.
10. W.K. Yeap, K.L. Holmes and M.E. Jefferies, What is an exit?, in *Proceedings of the Third International Conference on Automation, Robotics and Computer Vision*, 1994, 369-373.
11. S.P. Engelson, *Passive map learning and visual place recognition* PhD Thesis, (Yale University, 1994).
12. D.M. Kortenkamp, *Cognitive Maps for Mobile Robots: A Repesentation for Mapping and Navigation* PhD Thesis, (University of Michigan, 1993).
13. K.J. Holyoak and W.A. Mah, Cognitive reference points in judgements of symbolic magnitude, *Cognitive Psychology* 14, 1982, 328-352.
14. S.C. Hirtle and J. Jonides, Evidence of hierarchies in cognitive maps, *Memory & Cognition* 13(3), 1985, 208-217.
15. E.K. Sadalla, W.J. Burroughs and L.J. Staplin, Reference points in spatial cognition, *Journal of Experimental Psychology: Human Learning & Memory* 6(5), 1980, 516-528.
16. B. Tversky, Distortions in cognitive maps, *Geoforum* 23(2), 1992, 131-138.