Multiple Resolution Mapping for Efficient Mobile Robot Navigation

Kris de Meyer, Oliver Lemon, and Ulrich Nehmzow
Department of Computer Science
University of Manchester
Manchester M13 9PL
United Kingdom
lemonoj@cs.man.ac.uk, u.nehmzow@cs.man.ac.uk

5.9.97

Keywords: multiple resolution maps, robot navigation, topological mapping, computational efficiency

Abstract

We describe a novel approach to topological mapping for mobile robotics, and its advantages over existing proposals. The central idea is to develop a hierarchical "multiple resolution" representation from the robot's perceptions of its environment, computing only at a relevant level of detail for a given navigation task. The robot may then perform the task of *efficiently* navigating to a specified location. The resulting system is efficient since computation occurs only over the simplest representations necessary for each stage of the navigation. "Zooming" to a higher resolution of representation is map-driven, and the different resolutions are achieved by a uniform method using very small artificial neural networks of 4 nodes. Thus, as well as allowing efficient navigation, the robot's development of its map also uses few computational resources.

1 Motivation

In concert with the topological mapping community in mobile robotics (eg: [KB93, DFH93]), we believe that a strategy for mapping a robot's environment should not, in the first instance, rely on information about odometry. Such dead reckoning navigation systems rely on proprioception (perception of internal states), and are subject to incorrigible errors due to wheel slippage. Instead, we believe that representations must be anchored in "perceptual landmarks", obtained through exteroception (perception of the external environment). Such perceptual landmarks will represent those features of the environment which are salient with respect to the robot's own perceptual apparatus. The concept of "multiple resolution" in mapping an environment seems a natural choice to us, and is likely to result in increased efficiency of navigation (both with regard to time and energy resources). Multiple resolution mapping facilitates the allocation of *minimum* resources for the task at hand. To increase resolution (and thus demands on computation and memory) is only required under specific circumstances, such as the approach of a goal location.

Proc. TIMR UK 97, British Conference on Autonomous Mobile Robotics and "Scientific Methods in Mobile Robotics", Manchester 1997. Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 - 6161. Report number UMCS-97-9-1. http://www.cs.man.ac.uk/csonly/cstechrep/titles97.html

Many researchers (most notably [Lee96]) have asked what constitutes a "good" robot map of an environment. This issue of map quality is, of course, relative to particular tasks (see eg: [LP96]). We believe that details of computational efficiency of using topological maps have largely been neglected in appraisals of their suitability[1]. Besides computational efficiency, we consider the following criteria relevant to mapbuilding:

1. Spatial consistency: the map should faithfully represent the topology of some portion of the robot's environment. (i.e. the description provided by the map should be consistent with the actual environment.)
2. Coverage: the map should represent as much of the environment as is required for completion of the task.
3. Relevant detail: the map should only represent the environment in as much detail as is required for completion of the task.
4. Perception-based ontology: the map should only be constructed on the basis of primitives which the robot is able to perceive reliably.

Various proposals in the literature fail to meet at least one of the above criteria. As far as we know, none of them meet our requirement on relevant detail. For instance, in a long uniform corridor which the robot must reach the end of, it simply is not necessary to represent and compute over many separate locations along that path. More efficient journeys can be made if more detailed topology is only computed around target locations. Some theoretical results on topological mapping (eg: [DFH93]) fail to meet the last criterion, by relying on information that is difficult, if not impossible, to obtain from a real robot (for example the assumption that there is an enumeration of all path directions leading from each location). However, the results presented below are obtained from experimental data.

Our approach thus has the following features:

1. Topological mapping.
2. Multiple resolution - relevant levels of detail to task.
3. Computational efficiency of representation.
4. Perception-based representation.
5. Use of data from a real robot.

1.1 Related work

As far as we are aware, no other multiple-resolution approach to topological mapping exists. However, "hierarchical" representations of space have been proposed before, although in a quite different sense to that which we suggest. For instance Kuipers [Kui78] advocates the use of a hierarchy of spatial representations, moving from topological to geometrical information, but this hierarchy consists of different classes of representation, rather than of different levels of resolution within the same representation scheme.

Various proposals also exist with regard to the use of self-organising feature maps (SOFM) [Koh90] in mapping strategies. Typically, sonar and infrared readings are fed into (computationally costly) high-dimensional (e.g. 10x10) SOFMs in order to develop a single representation of the environment. In contrast, our approach uses just three linear 4 by 1 SOFMs to map feature space at three levels of detail. Tree-structured feature mappings are employed by [KO90] (although they are not maps in our sense), and are known to facilitate search algorithms of low computational complexity.

[1] Of course, it is well understood that topological maps are more efficient than metric ones.

2 The "Efficient Delivery" Task

The particular task that we have specified is that the robot must;

1. be able to develop a multiple-resolution perception-based topological map over a canonical path in its environment, and

2. use this ability to navigate to a specified location on the path, using higher resolution mapping only to disambiguate locations or when the target location might lie within its current region.

3 Benefits of multiple resolution mapping

As mentioned above, the computational cost of developing each level of representation is small (using only 4-node SOFMs) in comparison to other approaches. Further, the computational cost of planning routes over a small number of locations, in the low-resolution representation, is clearly less than that of the same problem over a more detailed representation. One might expect that this benefit could be obscured by the cost of developing a hierarchical representation, but since the more complex representations are developed only when required in order to complete the task (i.e. when the robot is in a target area), it is only in quite "uninteresting" environments (where perceptual aliasing is common, so that there are many target areas to be explored and disambiguated at higher resolutions) that computational costs might outweigh benefits. Thus, navigation should be faster, and more energy efficient, than in existing topological navigation methods. Such comparative studies are for future research at the present time.

4 Experimental set-up

4.1 The robot

The robot used for these experiments was a Nomad 200 mobile robot (see figure 1), named (after its serial number) "FortyTwo".

Figure 1: The Nomad 200 mobile robot "FortyTwo"

This hexagonal robot is equipped with sixteen ultrasonic range-finding sensors (range up to 6.5 m), sixteen infrared (IR) sensors (range up to 60cm), twenty tactile sensors and a monochrome CCD camera. The robot is driven by three independent AC motors for translation, steering and turret rotation. A 486 PC is the main controller with several slave processors (Motorola 68HC11) handling the robot's sensors.

4.2 Experimental Procedure

First, FortyTwo was left to follow the wall of the environment that was to be mapped, using a wall following behaviour which had been acquired through unsupervised learning earlier, using an artificial neural network to associate perceptions with actions ([Neh96]). Sensor signals of all sixteen sonar sensors, all sixteen infrared sensors and odometric information were logged during this exploration. The robot therefore mapped its environment as perceived along one canonical path, but it should be noted that multiple resolution mapping is independent of canonical paths or particular exploration strategies.

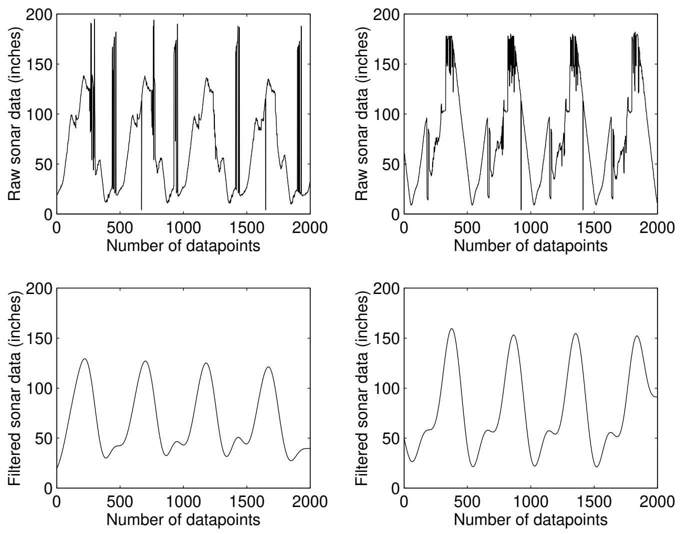

Figure 2: Raw and filtered sonar data for two different sonar sensors, before and after smoothing using a low pass filter. Spurious sonar signals, for example at datapoints 700 and 1700 (left) or datapoints 900 and 1400 (right) are removed, as is the high frequency content of the sensor signature.

Data from two different environments (a lab and an untreated corridor), for two independent experiments was obtained in this manner, and subsequently used to construct the multi resolution mapping off line. Through this procedure we ensure that results are relevant to robot mapbuilding (because "live" data obtained from a robot, interacting with its environment, was used), but also that a comparison between the different levels of resolution is possible, because in each experiment all three maps are trained using identical data.

4.3 Preprocessing of data

It is well known that sonar sensor readings are subject to noise. For example, "specular reflections" occur when the sonar pulse hits a smooth surface at too shallow an angle to be reflected directly. The reflection of the sonar burst by some other object in the environment generates a reading that indicates a larger free space than is actually available. Likewise, "crosstalk" (receiving another sensor's sonar burst) can indicate shorter distances to the nearest object than actually present. In figure 2 (left) such freak signals are visible, for example at datapoints 700 and 1200, which should exhibit identical sonar readings.

In order to remove of this noise, we filtered the sonar readings with a digital fourth order Butterworth filter with cut-off frequency 0.01 before starting the mapbuilding. The results for two arbitrary sensors are shown in figure 2. As can be seen, the noise has been successfully removed, without throwing away too much of the information content of the signal. The filtering operation causes an end-transient effect in the signal (see figure 2, bottom right, datapoint 2000), which can pose problems for the mapping mechanism, which we eliminated by removing the last 20 datapoints.

4.4 The mapping mechanism

The concepts of topological and multiresolutional maps, and their advantages, have been explained. We now focus on our particular approach. Self-organising feature maps (SOFM) [Koh90] have previously been used for mapping strategies (e.g. [NS91, Kur96]): sonar and infrared readings were fed into high-dimensional (e.g. 10x10) SOFM's, requiring a constant high level of computation. Here, we propose three levels of resolution, on each level we use a linear 4 by 1 SOFM to map feature space, thus providing a possible $4^3 = 64$ different place signatures. (A comparable single resolution mechanism would have to compute over 8×8 nodes). Each increase in resolution may multiply the number of recognised regions by 4, so the size of a complete high resolution map of an area is bounded by that of the initial low resolution representation. Whenever necessary, we can switch from a lower resolution to a higher one, always increasing the information content of the input signal. The computational cost can therefore be

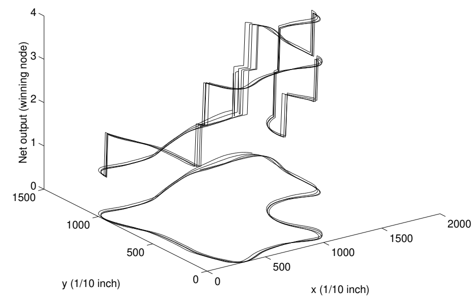

Figure 3: Coarse resolution mapping of the entire experimental area. The robot's path through the environment is plotted "underneath" the signature map, on the x-y axes. The eventual goal region (specified in a later stage of the experiment) is near location (0,800), the perceptual signature obtained there is "2".

reduced a large amount in certain "uninteresting" regions. Target regions, which the robot is to explore to the highest resolution, are specified externally to the system.

For the neural net training we used predefined "Matlab" functions from the neural net toolbox. The size of the linear (i.e. one-dimensional) self organising feature map was four units, neighbourhood size decreased from three to one over the training period. The learning rate decreased exponentially from 1 to (almost) 0. The SOFMs were trained over 5000 epochs (although fewer would almost certainly have sufficed), and each epoch consisted of an input chosen randomly from the data set.

On the resolutionally lowest level, the *whole* of the explored space is mapped into distinct regions (see figure 3). As input to the SOFM we use four infrared sensors only, spaced 90^o apart. The trained SOFM (training is completed after 5000 learning steps) then responds with one of four responses to perceptual stimuli received anywhere along the path.

A goal region is then specified externally, and at places that produce the same (low-resolution) output as this goal place a second network is trained at middle-resolution level (e.g. if the specified goal produces a "2" at low resolution, sensor signals perceived at all locations generating "2" at the low-resolution level are used to train the middle-resolution network). At the middle resolution level, four sonar sensors are used to train the SOFM, again spaced 90^0 apart. Regions that have been developed in this middle level now each have a signature like for example "2/1" or "2/4".

To differentiate further within those regions, a third 4x1 SOFM is trained, now using all 16 infrared sensor signals and all 16 sonar sensor signals of the robot.

The three levels of mapping resolution are thus developed by:

- Low level: 4 infrared sensors
- Medium Level: 4 sonar sensors
- High level: 16 sonar and 16 infra-red sensors

Through this process a tree-like topological map is constructed, where each low level region can be subdivided into smaller regions at a higher level of resolution.

5 Results

5.1 Experiment 1

In the first experiment a path of 13.5m was traversed 4 times, and 2000 sensor readings (sonar, infrared and odometry) obtained in a laboratory environment containing brick walls and cloth covered screens were used to produce the multi resolution mapping of the environment. Figure 3 shows the mapping of the entire environment at the lowest resolution level.

The "goal" location of the robot was externally specified as a location near the coordinates (0,800). At the lowest resolution level, the perceptual signature is "2", therefore all regions having this perceptual signature were subsequently mapped at medium resolution level (see figure 4). The goal region then

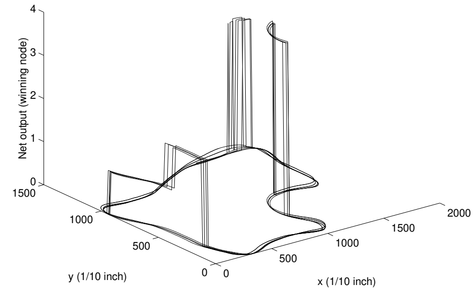

Figure 4: Medium resolution mapping of all regions giving response “2” at the low resolution level. The goal region now has a perceptual signature of “2/1” and is, in this case, already uniquely identified by this signature.

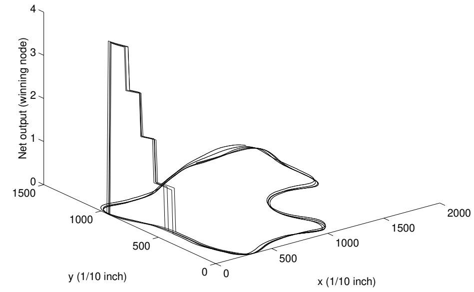

Figure 5: A further increase in resolution shows that the region “2/1” can be subdivided into four further distinct regions, pinpointing the goal location even more accurately.

has the perceptual signature “2/1”, and is, in this environment, already uniquely identified. However, a further increase in resolution shows that the goal region can be specified even more accurately (see figure 5).

5.2 Experiment 2

In a second experiment, FortyTwo was used to explore an untreated environment outside the laboratory, consisting of corridors, open spaces, doors, and irregularly distributed items of furniture. In total, a path of 35m was traversed once, and 2000 datapoints were logged and used for subsequent multi-resolution mapping.

Figure 6 shows the perceptual signatures at the lowest resolution level, obtained along the robot's path in this environment. The goal region was specified to be at distance 360 inches from the beginning of the path. The perceptual signature at the lowest resolution level is “4”, so that all regions having that signature were subsequently mapped at a medium resolution level (see figure 7).

The perceptual signature of the goal location now becomes “4/3”, so that all regions exhibiting that signature were then mapped at the highest resolution level (see figure 8). As can be seen from figure 8, the goal location is uniquely identified by the perceptual signature “4/3/3”

6 Summary and Conclusion

We have raised the issue of minimal resource use in the development of perception-based topological maps. Our criterion of “relevant detail” and thus computational efficiency of such a mapping strategy is one which, to the best of our knowledge, has not been noted before.

In order to decrease unwarranted processing load, the robot develops a hierarchical representation of its environment at multiple resolutions, so that space is partitioned coarsely where only a simple

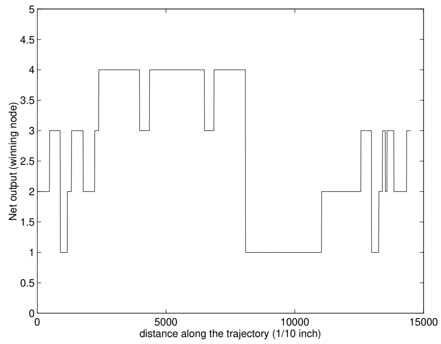

Figure 6: Coarse resolution mapping of the second environment in its entirety.

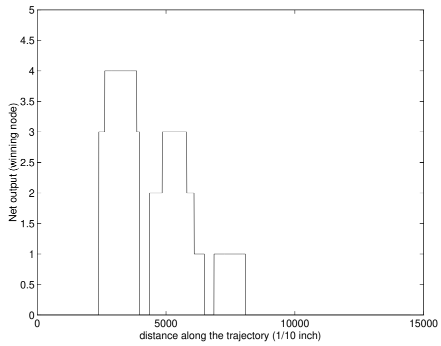

Figure 7: Medium resolution mapping of all regions of the second environment that have a perceptual signature of "4" at the lowest level.

representation is required, and the representation becomes more detailed ("zooms in") where the task at hand requires it.

The results presented above show that effective multiple resolution maps can be built from data obtained from a real robot, using a uniform mapping method. We believe that the main payoff from this novel representation scheme is yet to come, in the form of efficient navigation strategies. Efficiency arises not only in the development of the multiple-resolution map, but also in map-based navigation.

6.1 Applications of multiple resolution mapping

In this paper we have only been able to present the mechanism for developing multiple resolution topological maps. In future work we intend use the hierarchical representations for efficient navigation, in the following way.

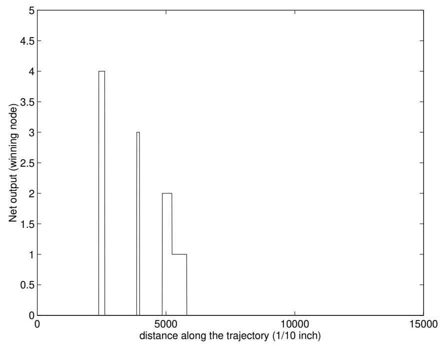

Figure 8: High resolution mapping of those regions of the second environment that have a perceptual signature of "4/3" at the preceding two levels.

6.1.1 The "zooming" strategy for multiple resolution navigation.

Target places may be described using a triple of signatures. The robot may then travel along a path, using only low resolution perception, until it reaches a region satisfying the first part of the triple. Having achieved this coarsely defined state, the robot increases sensor resolution ("zooms in"), and travels within the first region until it reaches a region satisfying the second label of the triple, if there is one. If no such region is found, the robot may switch back to low resolution perception ("zooming out") and travel until another region satisfying the first part of the triple is found. This process may be repeated until the target region is reached. Of course, possible perceptual aliasing may result in multiple solutions – but there are strategies for disambiguating such signatures ([DFH93, DN97]) on the basis of hypothesis testing or the use of local odometry, which can be adopted in order to solve this problem. We also intend to determine the efficiency gain of such a representation strategy, in comparison to other mapping and navigation techniques, through further experiments.

Acknowledgements

Kris de Meyer is an ERASMUS exchange student at the University of Manchester from ESAT Department, University of Leuven. Oliver Lemon is supported by the Leverhulme Trust (grant no. F/120/AQ).

References

[DN97] Tom Duckett and Ulrich Nehmzow, *Experiments in Evidence Based Localisation for a Mobile Robot*, In Proc. AISB workshop on "Spatial Reasoning in Animals and Robots", Manchester. Department of Computer Science, Technical Report Series, ISSN 1361-6153. Report UMCS-97-4-1, 1997

[DFH93] Gregory Dudek, Paul Freedman, and Souad Hadjres. Using local information in a non-local way for mapping graph-like worlds. In *13th International Joint Conference on Artificial Intelligence (IJCAI)*. Morgan Kaufmann, 1993.

[Koh90] Kohonen, T. The Self-organizing Map, In Proc. IEEE 78(9): 1464-1480, 1990

[KB93] B. Kuipers and Y. Byun. A Robot Exploration and Mapping Strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47 – 63, 1993.

[KO90] Pasi Koikalainen and Erikki Oja. Self-Organising Hierarchical Feature Maps", IJCNN International Joint Conference and Neural Networks, volume II, 279 - 284, IEEE, 1990

[Kui78] B. Kuipers. Modeling Spatial Knowledge. *Cognitive Science*, 2(2), 1978.

[Kur96] Andreas Kurz. Constructing Maps for Mobile Robot Navigation Based on Ultrasonic Range Data. *IEEE trans. on Systems, Man, and Cybernetics - Part B: Cybernetics*, 26(2):233 – 242, 1996.

[Lee96] David Lee. *The Map-Building and Exploration Strategies of a Simple, Sonar-Equipped, Mobile Robot; an experiments, quantitative evaluation.* PhD Dissertation, UCL, 1995.

[LP96] Oliver Lemon and Ian Pratt. Putting Channels on the Map: verisimilitude and spatial constraints in a semantics of Geographical Information Systems. In Larry Moss, editor, *Logic, Language, and Computation, volume 2.* CSLI Publications, Stanford (to appear).

[Neh96] Ulrich Nehmzow, *Flexible Control of Mobile Robots through Autonomous Competence Acquisition.* Measurement and Control, Vol. 28, No. 2. March 1995, pp. 48-54.

[NS91] Ulrich Nehmzow and Tim Smithers, *Mapbuilding using Self-Organising Networks*, in: Jean-Arcady Meyer and Stewart Wilson (eds.), *From Animals to Animats*, MIT Press, Cambridge Mass. and London, England, 1991, pp. 152-159.

[TB96] Sebastian Thrun and Arno Bucken. Integrating Grid-Based and Topological Maps for Mobile Robot Navigation. In *13th National Conference on Artificial Intelligence*. AAAI, 1996.

[YB96] Brian Yamauchi and Randall Beer. Spatial Learning for Navigation in Dynamic Environments. *IEEE trans. on Systems, Man, and Cybernetics - Part B: Cybernetics*, 26(3):496 – 505, 1996.