A Robust Modular Multi-Representational Mapping Model.

Simon J. EGERTON, Vic CALLAGHAN, Paul CHERNETT
Department of Computer Science, University Of Essex, Wivenhoe Park, Colchester, C04 3SQ, England
egers@essex.ac.uk vic@essex.ac.uk cherp@essex.ac.uk

26.3.99[1]

Abstract: This paper presents a modular mapping model that allows an autonomous mobile robot with inaccurate sensors to construct a map of an unknown, unstructured and unmodified environment. The robot constructs the map in real-time autonomously with no a priori knowledge. To achieve this topological and geometrical methods are brought together to produce a mapping model that is robust, efficient and modular with a flexible method for sensor fusion. The produced map being adequate for path planning and localisation tasks. Moreover, it is believed this method achieves these tasks more elegantly than other existing methods. Experiments are currently being conducted on a mobile indoor robot. The experiments will later be extended to outdoor robots. The inspiration for this model was found in the biological literature.

Keywords: Robotics, Perception, Neural Networks, Geometric Reasoning, Environment Modelling

1. Introduction

We believe that an autonomous mobile robot should have the ability to adapt to its target environment. This adaptability should extend as much as possible throughout the robots functionality. This adaptability increases the usefulness of the robot and its autonomy. Target environments that are unknown, unstructured and unmodified present an interesting challenge to the autonomous mobile robot. It must have a flexible functionality to deal with these three environmental 'U' factors. Part of the robot's functionality, particularly in non-reactive robots, may be to gather knowledge of its target environment as it explores. Commonly this knowledge is organised into some type of environmental representation, for example an environmental map. This mapping process should be flexible and adaptable, the robot constructing the map from scratch, with no a priori map or knowledge of the target environment. With these conditions the mapping function of the robot should only use the sensory impressions of the exploring robot.

Many mapping strategies have been suggested. We should recognise that the intended application and target environment of the robot will influence the type of mapping strategy used. In general mapping methods may be broadly categorised into three categories. The first, are those methods which are the "traditional" geometric or qualitative methods, such as (Elfes, 1989; Pagag and Nebot, 1995; Noborio et al, 1990). These methods are based on the accumulation of accurate geometric information about the world. This dependency on accurate sensor information makes these methods impractical for robots situated in the above types of target environments. The second category of methods is the topological, or qualitative, methods, such as (Zimmer, 1995; Kuipers and Byun, 1988). Rather than using rigid geometric information these types of methods use the robots sensory impressions more directly to construct an environment map. The third category of methods is formed from using both geometric and topological methods. In these methods the quantitative and qualitative information is clearly separated. Good examples of this hybrid approach are (Kurz, 1996; Thrun 1998).

The modelling method we introduce in this paper belongs to the third category of methods. The model was originally inspired from the biological literature (Cheng and Gallistel, 1983; Cheng, 1986). The modelling method aims for the robot to autonomously produce, in real-time, a map that models the world with a higher degree of resolution, without overfitting the map with the mapping structure, with the robots target environments conforming to the three environmental 'U' factors. Moreover, the model also being used for localisation and path planning tasks, with the path planning algorithms exploiting the proposed models structure for computational efficiency.

In the rest of this paper we describe the robot used in the experiments and also suggest some benchmark tests to test this and other models with. We then describe the proposed model along with some results.

2. The Robot

Currently all the experiments are performed using an indoor laboratory robot. Figure 1 shows the robot. The robot is equipped with a number of sensors. An onboard electronic compass provides angular information. This is subject to errors, electronic equipment and metal objects in the laboratory interfering with the surrounding magnetic field. Two stepper motors drive the robot's wheels. Internal position information is derived from the number of steps taken by the stepper motors and the

[1] Proc. TIMR 99 "Towards Intelligent Mobile Robots", Bristol 1999.
Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 – 6161
Report number UMCS-99-3-1. http://www.cs.man.ac.uk/cstechrep/titles99.html

compass angle information. This crude method ensures the internal position information is never very accurate and accumulates error if unchecked. A measure of the mapping model's robustness is the amount of positional error tolerated before the map becomes unreliable. A set of touch sensors surrounds the robot perimeter and provides positive feedback when the robot bumps into objects. The robot has a set of eight simple ultrasound sensors distributed evenly around its perimeter. These sensors provide range distances relative to the robots current position. The readings are subject to the typical noise associated with ultrasound measurements i.e. reflections from surfaces, environmental temperatures, surface textures. The robot has a vision sensor and this is described below. The robot has an onboard 68040 processor with 8MB of memory and runs the real time operating system VxWorks.

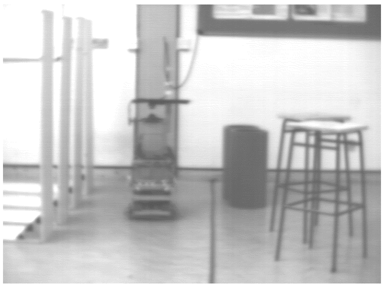

Figure 1 *Showing the lab robot situated in an example indoor test environment. Objects seen are wooden strips, chairs, cones, and a wall.*

3. Test Environments And Benchmarks

To help develop and evaluate the proposed model we require a set of controlled test environments and benchmarks. At present, there are no standard environments or benchmarks to evaluate existing models. We therefore have the freedom to specify such environments. Presently we have specified the first and second test environment and the final benchmark. Intervening test environments and benchmarks will be specified through the course of the project. All of the environments adhere to the three environmental 'U' factors. The first types of environment will be static. Dynamic environments will be considered at a later stage.

The first of the test environments is designed for development purposes only. It is an oversimplification of the typical environment the robot is expected to navigate through and is illustrated in figure 2a. This "straight line" environment simplifies the navigation process and provides a very controlled method for observing the effects different objects and parameters have on the model.

The second test environment is illustrated in figure 2b and is simply an extension of the first test environment into two dimensions. Its purpose is to test the stability of the models Perception Space component(s), which are explained below. The results from this test environment are expected to be independent of the exploration strategy used by the robot.

For the final test environment or "benchmark" the robot will start from a "home" position and explore the environment, constructing a map along its way. Once a sufficient map has been constructed the robot will be placed in a random position and instructed to find its way back to the home position. The success of the model will be a measure of the distance error away from the starting "home" position and the length of the path taken to get there. These notions are illustrated in figure 3. We believe the model presented here will compare very favourably with other models surveyed so far.

(a)

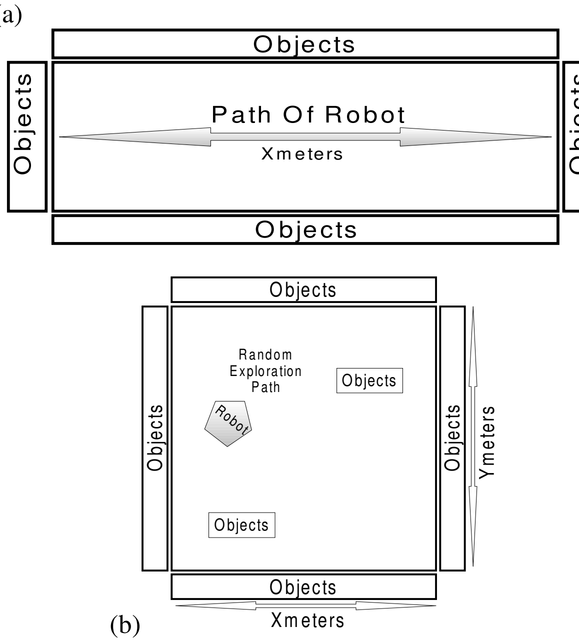

Figure 2 *(a) The first type of test environment, type one. The robot navigates along a straight line bounded by objects. The environment parameters are the length of path and the types of objects used. Typical objects for an indoor environment include for example wooden planks, cones, walls, and chairs. (b) The second type of test environment, type two. The environment is set up with an area bounded by a set of objects. The main parameter of interest here is the exploration path of the robot.*

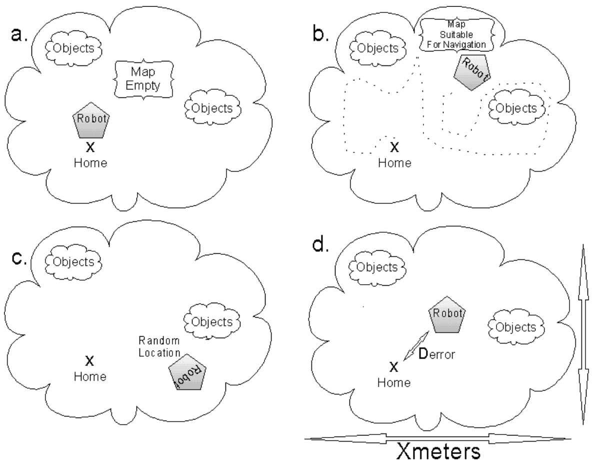

Figure 3 *The final benchmark: a) The robot starts in an unknown environment and explores. b) The robot has explored the environment and has constructed a map sufficient for navigation and localisation tasks. c) The robot is taken to a random place and instructed to find home. d) The robot reports it has found home, the measure of its success is the distance away from the home marker and the length of the path taken to get there.*

4. The Model

The inspiration for the model we present below was derived from an interpretation of the biological literature (Cheng and Gallistel, 1983; Cheng, 1986; Bennett, 1996; Gallistell and Cramer, 1996). Presented in this literature is experimental evidence that mammals compute and make use of geometric information in forming environmental maps (Cheng and Gallistell, 1983). Experimental evidence is also presented that topological information is also computed. These two types of information are brought together by a "Geometric Frame Module" and a set of "Feature Modules" (Cheng, 1986). The Geometric Frame module models the shape of the environment. The geometrical information to do this is derived from the sense, for example, optic flow calculations (Bennett, 1996). The non-geometric information is located on the Geometric Frame by addressing the appropriate Feature Module. For simplicity a Feature Module can be seen as detecting a specific non-geometric feature. It is also suggested that the complexity of the Feature Modules increase with the complexity of the animal, however the Geometric Frame Module remains the same. These notions are illustrated in figure 4. The model we now present is an interpretation of this biological model. To explain the model we introduce two concepts. Firstly the notion of a "Perception Space", this is analogous to the "Feature Modules". Secondly the notion of a "Geometric Space", this is analogous to the "Geometric Frame Module".

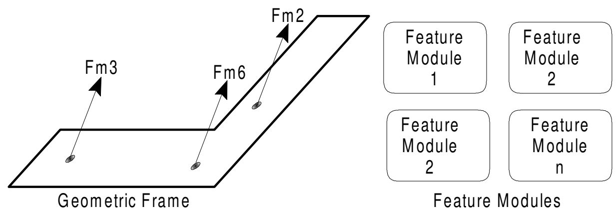

Figure 4 *Illustrating the Geometric Frame Module and the Feature Modules. The geometric frame is featureless, and addresses the feature modules. For example, feature module one may equate to food, feature module two may equate to a burrow, etc.*

4.1 The Overall View

The robot is equipped with a range of sensors through which it senses the world that surrounds it. Exploring the world the robot should build its map from scratch using only its sensory impressions. These sensory impressions are defined as sensor element activation's relating to some fixed physical location in the environment (generally not directly equatable to objects in the "human" perception domain) and we term these "Perception-Signatures". To illustrate this notion, let the universal set be the set of all perceptions perceivable by any sensory means, then individual sensors form subsets of this universal set. Each sensor can be seen as acting as a filter, only allowing through perceptions that are specific to it (Wilson, 1991). By this, each sensor has its own associated set of perceptions, this is visualised in figure 5.

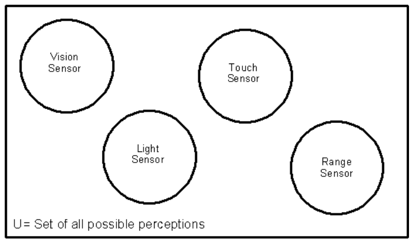

Figure 5 *Illustrating the notion of each sensor filtering the universe of perceptions so forming its own perception set.*

The sensors space of perceptions is further constrained by the type of environment the sensor is situated in, forming an "environmental subset". This environmental subset of perceptions, or "Perception-Signatures", is constrained further, forming another subset, which depends on how much of the environment the sensor gets to see. It is this subset which forms the "Perception Space". Within this space, the perception-signatures are categorised for their similarity relating to similar "Perception Areas" in the environment. For example, if the robot finds itself in a farm environment equipped with a camera then its environmental subset of perceptions are all things perceivable with the camera on the farm. However the robot spends its days trudging around the fields and its perception space only consists of things like fences, hedges, wooden poles, bales of hay etc. Therefore, as the robot explores the environment, perception-signatures are categorised and perception spaces formed for each type of sensor it has. The perception space starts as an empty set, growing as more of the environment is seen and should stabilise once the entire set of perception-signatures in environment have been seen. However, if new parts of the environment are discovered then the perception space should accommodate the potentially new set of perception-signatures.

The "Geometric Space" is a geometric framework and its purpose is to simply relate geometric areas to perception areas in the environment where they occur. This is needed since it is unlikely that a perception area will be unique to one area of the environment i.e. "Perceptual Aliasing" (Duckett and Nehmzow, 1997). Indeed, if every place in the world visited by the robot appears unique to its sensors, then the robot can use this information alone to build a navigable map.

The geometric space addresses the perception space to relate geometric areas to perceptual areas. A perception area may be referenced more than once by the geometric space. In our farm environment for example, there may be more than one bale of hay lying around in the field. As explained later in this article inexpensive path planning algorithms can be implemented with careful consideration to the representation of the geometric space. In addition, the robot should expect to see a certain sequence of perceptions along any planned path. This allows the robot to check its progress along the path and not be entirely dependent on its kinematics.

While the perception space should stabilise once the entire set of perception-signatures in the environment have been seen the geometric space may continue to expand as the robot explores more of the world. For example in our farm environment after exploring an entire field, the perception space will be stable. Now there may be a number of similar fields requiring the robot's attentions. Therefore, the geometric space expands to accommodate the new fields as the robot explores them, but the perception space remains unchanged since all the fields contain similar objects and appear similar. However, should the robot then go to explore the barn where the hay is stored the perception space should accommodate the potentially new set of perceptions. The combination of the perception space and the geometric space is illustrated in Figure 6.

The key to this model is the partitioning of the perception space into perception areas. It is not important what the perception areas relate to in the physical world, all that matters is that there are a set of areas in the environment which are distinguishable by the robots sensors. Since it is the ability to distinguish different perception areas that allow the geometric space to be constructed and allow path planning and localisation tasks to be performed. Using our farm environment for example, Pa1 could relate to open field, Pa4 could relate to bales of hay and Pa3 could relate to a fenced area, etc. On the other hand, they could relate to something far more abstract. Again, the importance lies with the robot being able to distinguish various perception areas in the environment. Whether or not the perception areas relate to objects "humans" can identify with does not matter, although of course this would be convenient. Perception areas are simply labelled by number, anything more meaningful requires the robot to have an "understanding" of what it perceives. Alternatively a human operator could observe the robot and try to match perception areas with some appropriate label

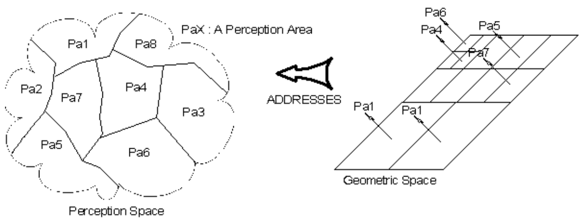

Figure 6 *The Geometric Space referencing a Perception Space. This combination provides navigation information and a context for the robots perceptions. The Perception Space shown has been partitioned into Perception Areas.*

4.2 Sensor Fusion, Localisation and Path Planning

A perception space is formed for each type of sensor the robot has and a geometric space is constructed for each perception space. This process can be achieved in parallel. Since each geometric space shares the same environment space, simply adding the geometric spaces together will fuse together the sensor information, thus forming a navigable map. Moreover, no information is ever lost, since each sensor has its own perception space and geometric space and the addition does not effect this. This method also provides additional flexibility. It is possible to select which perception and geometric information is used to provide the final navigable map, based on some criteria. This is useful, for example one sensor type may become damaged, or be giving false readings. In this case the sensors perception and geometric spaces may then be ignored and navigable map formed from the remaining good sensors. These notions are illustrated in figure 7.

If the robot becomes lost, for whatever reason, the robot can derive its location using the map it has built. Moreover, it can do this without referencing any positional information by making use of the perceptual space. When the robot is lost, by definition, it will be in any one of a set of possible locations. In this situation, the robot should wander around gathering evidence to find out where it is. The evidence is more robust if the robot has a set of sensors appropriate for the environment and combines the information together as described above. This way the robot has more evidence to reference locations with. The localisation procedure is formalised by the algorithm below. It relies on neighbourhood information in the combined geometric space to reduce the robots set of possible locations.

Location Algorithm:

1. Note the current perception area(s).
2. Create a set of locations from the combined geometric space that share the above perception area(s).
3. Initialise the localisation hypothesis set with the set produced in step 2.
4. While location hypothesis set contains more than one element
5. Wander around until the perception area(s) change.
6. Create a set of locations from the combined geometric space that share the above perception area(s) and neighbour a location in the current location hypothesis set.
7. Set the location hypothesis set with the locations created in step 6.
8. Location equals location in the hypothesis set.

Here the robot wanders around the environment until it finds a particular sequence of perceptions that can be identified uniquely in the map. For example, in our farm environment if there is only once place where the robot would pass a gate followed by a hedge then by a fence, the robot can locate itself on the map after this sequence of perceptions. Successful localisation depends on the environment being varied enough to allow the robot to reduce the set of possible locations to one location. We also suggest that while the robot is lost a separate "lost map" be created. This information may then be integrated with the main map once the robot has re-localised itself. In this way no information is wasted.

It should be noted the location algorithm does not account for all eventualities. For example, if the robot encounters a new perception that does not have a neighbour in the current hypothesis set? This is the result of any one of three events:

1. Firstly and the most critical, the map was incorrectly built. This should not be allowed to happen.

2. The robot has moved out of the mapped area. In this case, the robot really is lost.

3. The environment has changed. In this case, either one or both of the above facts could be true as well.

We are aware of these problems and do not have any solutions to them now. Currently with the model, the only

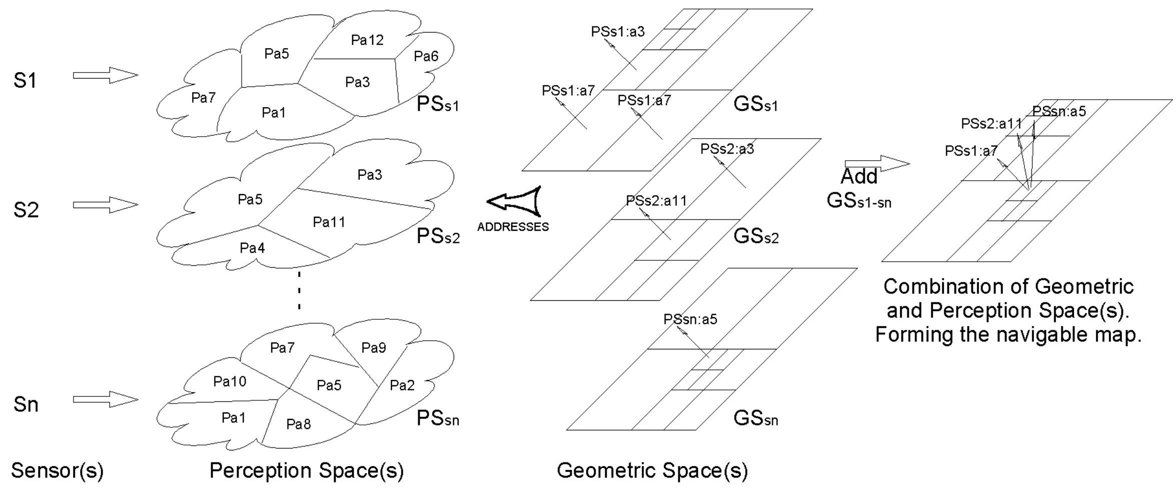

Figure 7 *Illustrating the notion of sensor fusion within the model. Sensor information may be combined through adding the geometric spaces together. This process can be selective of the sensors used based on some performance measure or other criteria.*

solution if this situation occurs is to restart the location process. However, are these location problems any different to those "humans" would have in similar situations? Would solving these problems result in a system out performing human capabilities?

The model is readily used for path planning. Paths are generated in the combined geometric space. The robot can then navigate along the planned path with a set of expectations. These expectation are what it should expect to "see along the way" in terms of perception areas. This feedback from the perception space can be used to check if the planned path is being followed. If the set of expected perceptions along the path is not met, then we may assume that either the robot is lost or that one of the above three conditions holds.

5. Implementation

The function of the perception space module is to partition a perception space into suitable perception areas. This partitioning function is a data classification problem. A "Growing Cell Structure" neural network (Fritzke, 1993) is used for this classification process. The growing cell network is a self-organising neural network and works on the same principles as the well-known Kohonen neural network (Kohonen, 1988). The structure of the Kohonen network is fixed. The structure of the Growing Cell network is not fixed and is able to evolve with the data, growing and shrinking to best categorise the data. These are desirable features since the robot will be exploring an environment that is of unknown size, shape and complexity. The cells of the network represent the perception areas and the connections between them express their topological similarity.

The geometric space module relates perception areas to geometric areas in the environment where they occur. The Geometric Space is modelled with a "Growing Quadtree" a modification of the standard Quadtree (Samet, 1984). The quadtree is based around the principle of recursive decomposition. Areas are viewed as quadrants, and quadrants are recursively decomposed until quadrants are either homogeneous or the minimum quadrant size has been reached. In our case, homogeneous areas represent areas with the same perception area within a particular sensors perception space. Hence uniform areas are mapped with the minimum of quadrants and the resolution of the represented area is defined by the minimum quadrant size. The quadtree representations of the geometric space can be combined simply by forming the union of the quadtrees. So forming the overall navigable map.

Quadtrees offer a dynamic and efficient method for representing areas, the resolution to which an area is represented can vary as well. In addition, tasks such as path planning can make use of the efficient algorithms developed for the Quadtree structure (Kambhampati and Davis, 1986; Zelinsky, 1992).

6. Initial Results

The Perception Space is the essential component of the proposed mapping model and initial experiments are concerned with its development. Indoor environment types are used.

The readings from the robot's ultrasound sensors are subject to noise i.e. changes in environmental temperature, surface reflections and signal absorption. We have found that pre-processing the raw ultrasound information with a simple moving median filter, as often used to de-speckle images (Jain et al, 1995) significantly reduces noise and filters out false readings. The effects of applying the moving median filter to a raw sensor profile are illustrated in figure 8.

An ultrasound perception-signature consists of eight readings, one reading taken from each of the robots sensors. The ultrasound sensors are evenly distributed around the robots perimeter to achieve a "Panoramic" view of the environment. Panoramic views are desirable since they simplify location recognition. However, the eight readings do not form a very high-resolution panoramic perception-signature. An example of a high-resolution ultrasound panoramic perception signature is illustrated in figure 9. Rotating the robot through 360^0 and sampling the same ultrasound sensor every 3^0 produced this. This is equivalent to 120 sensors evenly distributed around the robot's perimeter sampling the environment at the same time.

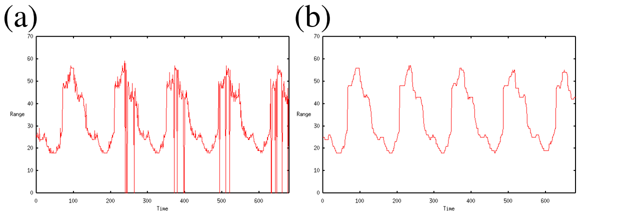

Figure 8 *Illustrating the Median filter process. Shows the readings from one sensor as the robot moves through a test environment (a) Shows the raw ultrasound data (b) Shows the median filtered data, filter width of seven.*

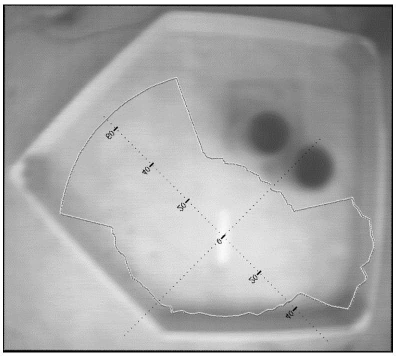

Figure 9 *Illustrating a high-resolution ultrasound panoramic perception-signature. Showing an overhead plan view of the test environment. The perception signature has been overlaid onto the test environment, the robot is centred at the origin.*

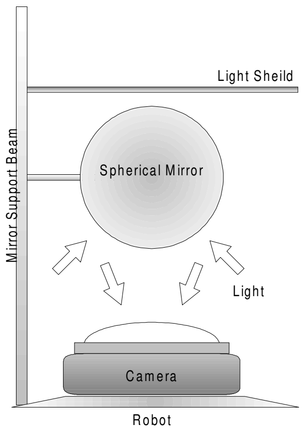

Figure 10 *The arrangement of the vision sensor on the robot. Ideally the mirror should be a Hyperbolic (*Svoboda et al 1998) *this prevents image compression as the panoramic image horizon is approached and also maintains geometrical relationships.*

Vision sensors have the potential of acquiring high-resolution panoramic images. We make use of an elegant idea (Rushant and Spacek, 1998) that allows this to be achieved simply and reliably. The basic arrangement of the vision sensor is illustrated in figure 10 and is situated on top of the robot. The robots electronic compass is used to compensate for the orientation of the sensor arrangement, which is relative to the robots current orientation.

The camera is a simple black and white CCD camera and is focused on to a spherical mirror. Light is reflected off the mirror into the camera producing a panoramic image. The sensor is in the early stages of development. A typical image from this arrangement is illustrated in figure 11a. These "raw" panoramic views are processed to form perception signatures.

The perception signature is produced by extracting light intensity information by segmenting the raw image into regions of m^o arc and taking the arithmetic mean of the intensities within the regions. This produces a n-dimensional vector that is the perception signature, where $n = 360^o / m^o$, the value of m^o determines the resolution of the perception signature. Averaging within each region may also be restricted between two radii P_{rr}. The value of P_{rr} will "focus" the perception signature on various parts of the environment. Low values focus on areas close to the robot and higher values focus on areas further away. The amount of "focusing" is dependent on the diameter of the spherical mirror. Larger diameters allowing more focusing control. These notions are illustrated by figure 11b.

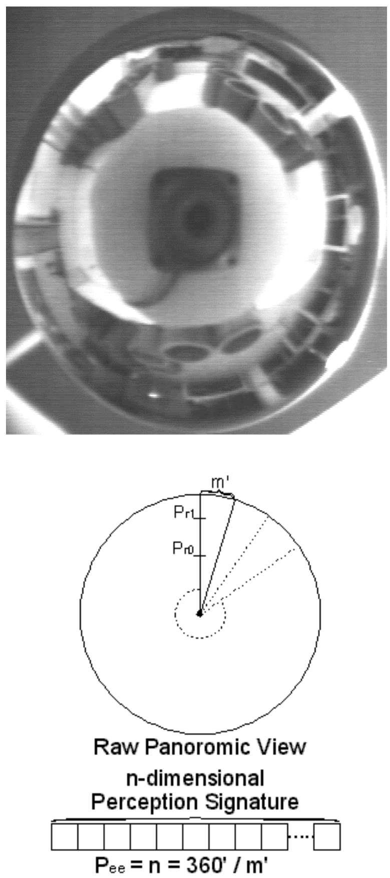

Figure 11 *a) Panoramic image of a location in a test environment, as captured from the robot. The camera is seen in the centre with the surrounding environment showing obstacle cylinders, walls and open space. b) Perception Signature elements are the arithmetic mean of corresponding region in the raw view. Element values may range from 0 to 255. Perception Signatures may be displayed as a grey scale image.*

The work of (Nehmzow and McGonigle, 1988) is an example were light intensity information is used successfully for navigation purposes. However, in this case explicit light intensity devices are used. In comparison, the vision sensor described here is more flexible. Light intensity vectors can be produced at varying degrees of resolution by changing a few parameters and without the need for additional sensor devices. In addition, this sensor can also be used for other vision-related functions.

For the mapping model to work, the sensors on the robot must suit the environment. If there is measurable variation within the environment, then a suitable sensor is one that is able to detect these variations. For example, consider an empty environment whose surface is titled, each tile having a unique colour. Clearly ultrasound sensors do not really suite this environment. However, a colour camera sensor might be just what is required. With the same notions we consider the spherical vision sensor described above. The vision sensor uses a grey scale CCD camera to detect light intensities. So the environments this sensor is suited to are ones in which there are variations in light intensity. This can be easily achieved within indoor environments by way of obstacles and artificial lighting. However, the weakness in using light intensity information directly is in that it is fragile to variations in surrounding light levels. It is ideally suited to indoor environments where the light levels can be controlled. We see the vision sensor as a multi-resolution experimental tool to enable the notions of the mapping model to be experimented with. It provides a great deal of flexibility with minimal computational and hardware costs. If we are to use this vision sensor in an outdoors environment, then we suggest that a colour camera be used instead. The RGB component values could be used in the same way the light intensity information is used for the indoor environments. Each component forming separate perception signatures each associated with their own perception space.

We now described a few results obtained with the vision sensor. The robot was let to wander once through a type one environment as described above. This trains the Perception Space. The Perception Signatures are processed in real-time. They are used as input vectors to the Growing Cell structure where they are categorised. Each category is a Perception Area. A set of Perception Signatures is illustrated in figure 12. After training, the robot was let to wander through the environment again and the recognised Perception Areas recorded. We expected the Perception Space to produce stable Perception Areas. Figure 13 shows the results of two test runs. Each test run was performed with different environment settings. It is clearly seen that locations in the environment have successfully been classified and then later recognised. These results are very encouraging and show the Perception Space and vision sensor are working as expected

7. Further Work

The results obtained to date have been very encouraging. Regarding the implementation of the Perception Space. We need a method that categorises the Perception Signatures into Perception Areas. There are many algorithms to do this. We chose the Growing Cell structure, but now we believe this to be too general for its purpose. In general, it is designed for mapping unknown probability distributions of unknown shape and size, also allowing dimensionality reduction in the input data for visualisation purposes. Furthermore, it has many parameters that effect its behaviour. The Perception Signatures are not random, but more continuous. For example, considering a static environment. As the robot moves through the environment, the Euclidean distance between consecutive perceptions will be small, so long as the sampling frequency of the sensor is high relatively to the speed of the robot. Generally the current perception will be similar to the previous perception. Large changes in Euclidean distance will signify a significant change in the environment. This “temporal” element of the Perception Signatures is ignored by the Growing Cell structure, but it could be used to simplify the categorisation algorithm of the Perception Space. We are investigating the use of the “Box” classifier (Jeffery and Picton, 1995) as an alternative. It has fewer parameters and is able to make use of the temporal relationship between successive Perception Signatures. These temporal notions are illustrated in figure 14.

Considering the whole model. The model needs relative co-ordinate information to construct the geometric space(s). Currently co-ordinate information is derived through the robot’s stepper motors and electronic compass. This method is error prone and accumulative. A better method would be to calculate egocentric co-ordinate information through the robots external sensors. For example, with the vision sensor, this might be done for instance by some optical flow field method (Svoboda et al 1998). We are currently investigating these possibilities.

8. Conclusion

This paper has presented a biologically inspired mapping model. It allows a mobile robot to model its environment. It can do this autonomously in unknown, unstructured and unmodified environments with no a priori map or knowledge. Its modular multi-representational approach allows new sensors to be added simply. Moreover, sensor fusion is a simple matter of selecting the appropriate sensor perception spaces and adding the respective geometric space(s) so forming the navigable map. The model is readily used for self-localisation and navigation tasks. The structure of the model exploiting proved efficient navigation algorithms. Currently there are no standard benchmark environments to allow the evaluation of this and other mapping models. We are developing a set of such benchmarks. The final benchmark as described above is designed to challenge this and existing mapping models. Moreover, we believe the model presented here will compare very favourably with other models surveyed so far.

The Perception Space is the critical component of the model. Initial experiments are concerned with developing it into a robust module. The results so far are very encouraging and supportive of the proposed model. Future experiments will include outdoor robots situated in equivalent outdoor environments.

References:

Cheng, K. 1986. A Purely Geometric Module in the Rats Spatial Representation, Cognition 1986.

Cheng, K., Gallistel, C. R. 1983. Testing the Geometric Power of an Animal's Spatial Representation, Animal Cognition

Duckett, T., Nehmzow, U. 1997. Experiments in Evidence Based Localisation for a Mobile Robot, Proc. AISB workshop, Manchester 1997.

Elfes, A. 1989. Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation, PhD thesis, Carnegie Mellon University, 1989.

Fritzke, B. 1993. Growing Cell Structures - A Self Organising Network for Unsupervised and Supervised Learning, Technical Report 93-026, International Computer Science Institute, Berkely, California, May 1993.

Noborio, H., Kondo, K., Noda, A. 1990, IROS90

Jain R., Kasturi R., Schunck, B. 1995. Machine Vision, McGraw-Hill International Books, 1995.

Jeffery, J., Picton, P. 1995. Mechatronics, Volume 2, p37.

Kambhampati, S., Davis, L. 1986. Multiresolution Path Planning for Mobile Robots, IEEE Journal of Robotics And Automation, Vol2 No.3, September 1986.

Kohonen, T (1988), Self organisation and associative memory, New York, Springer Verlag

Kurz, A. 1996. Constructing Maps for Mobile Robot Navigation Based on Ultrasonic Range Data, IEEE Transactions on Systems, Man, Cyber. April 1996, p223-242.

Pagac, D., Nebot, E.M. 1995 Global and Local Map Building and Maintenance for Autonomous Mobile Vehicles, Proc. National Conf. of the Australian Robot Association, Melb. 1995

Nehmzow, U., McGonigle, B. 1988. Robot Navigation by Light, European Conference on Artificial Life, May 1993.

Rushant, K., Spacek, L. 1998. An Autonomous Vehicle Navigation System using Panoramic Machine Vision Techniques, Proceedings of the International Symposium on Intelligent Robotics Systems, ISIRS98, India, Jan. 1998.

Samet, H. 1984. The Quadtree and Related Hierarchical Data Structures, Technical Report 23, Maryland University.

Wilson, S. 1991. The Animat Path to AI, SAB91.

Svoboda, T., Pajdla, T., Hlavac, V., Motion Estimation Using Central Panoramic Cameras, In IEEE Conference in Intelligent Vehicles, Stutgart, Germany, October, 1998.

Thrun, S., Bucken, A. 1996 Integrating Grid-Based and Topological Maps for Mobile Robot Navigation, AAAI96, August 4 - 8, 1996, Portland, Oregon

Yagi, Y., Kawato, S. 1990. Panoramic Scene Analysis with Conic Projection, IROS90

Zelinsky, A. 1992. Mobile Robot Exploration Algorithm, IEEE Robotics and Automation. Vol8 no6 Dec. 1992

Zimmer, U. 1995. Self-Localisation in Dynamic Environments, IEEE/SOFT Int. Workshop, BIES'95, May 30-31, 1995 Japan.

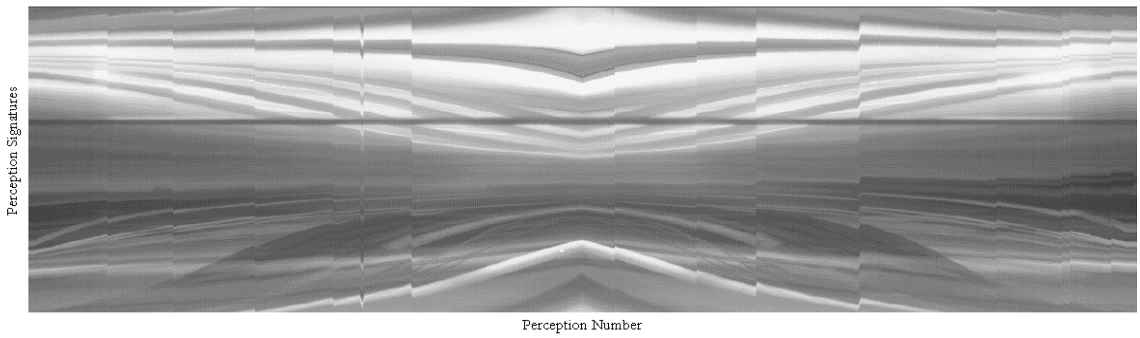

Figure 12 *The processed Perception Signatures from a test run through a type one environment of 5 meters in length. For display purposes the signature elements are displayed in sequence forming a continuous grey level image. The robot wanders forwards through the environment and after reaching the end wall wanders back again. The perception signatures above reflect this symmetrical path. The solid strip seen across the signatures represents the spherical mirrors fixing mount. Vision sensor parameters* P_e *= 360,* P_{rr} *= [83..103].*

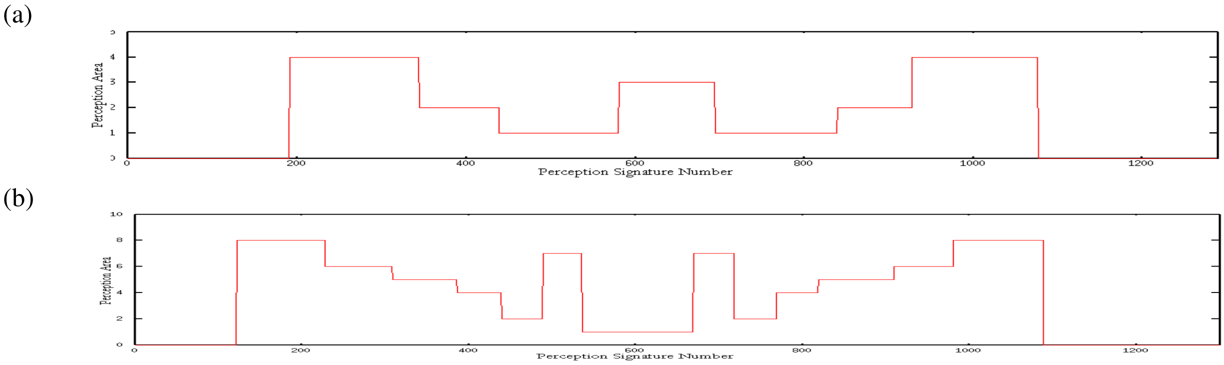

Figure 13 *Illustrating some Perception Space results; Growing Cell structure parameters* ε_{init} *= 0.15,* ε_{nei} *= 0.035,* $\Delta\varepsilon_{bmu}$ *= 0.01,* $d_{accuracy}$ *= 0.10,* B_{size} *= 500,* B_l *= 5,* N *= 1; Vision sensor parameters* P_e *= 360,* P_{rr} *= [83..103]. (a) Shows the Perception Areas created within an environment similar to that show in figure11. The environment used in (a) was modified to include more objects and the results of the categorisation are shown in (b). As expected they are different. Both sets of recognised Perception Areas are symmetrical as the robot moves forwards through the environment and then backward.*

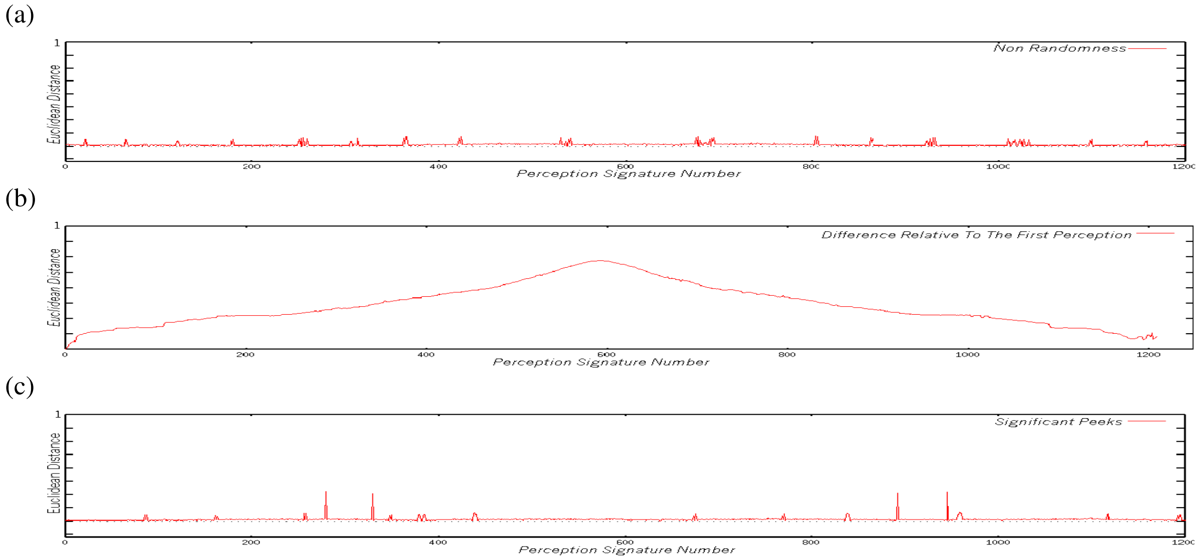

Figure 14 *Illustrating the temporal relationships between consecutive perception signatures from the vision sensor. The robot navigates through some symmetrical environments. (a) Euclidean differences between consecutive signatures. The small ripples are rotational errors in the raw perception image processing (b) Euclidean differences between the first signature and successive signatures. (c) Same as (a) the sharper peeks indicate an abrupt feature, in this case passing by a brightened corridor.*