Learning an approximate map of the environment by unsupervised bimodal landmark exploration[1]

Lambert Schomaker [a] Rudolf Fehrmann [a]

[a] AI Inst./University of Groningen, Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands

Abstract

Technological methods for navigation have achieved a high level of perfection, currently. However, this perfection is mostly achieved at the cost of additional artificial equipment which is extraneous to a freely-navigating agent such as an autonomous robot. Cognitive, biological navigation is based on functionality which is intrinsic to the cognitive system and thus more flexible and autonomous in the true sense. In this paper, exploratory experiments in navigation are performed, based on proximity events, visual landmarks, and distances traveled. It will be shown that a robot is able to learn a Kohonen self-organized landmark map of image and sonar data. An approximate 2-D map of the environment can be computed on the basis of the two major principal components within a sparse distance matrix between landmarks.

1 Introduction

Technological systems for navigation, i.e., satellite-based global positioning (GPS) and laser-based triangulation have achieved a level of utmost perfection. Also, artificial sensing such as 2-D laser range finding may provide stable information for navigational mapping [8]. However, it is amazing how good animals are at navigating through their habitat, without using additional hardware, laser, absolute coordinate systems or rectangular-grid maps. On the contrary, in biological systems, ego-centric and landmark-based coding play a predominant role, yielding navigation behavior which statistically reveals an underlying robust and flexible mechanism which supports survival in a real environment. Therefore, it remains a scientific challenge to develop models for map learning which are based on biologically plausible representations and learning schemes. The current work is intended to be a basic step in the direction of robust, cognitive navigation modeling. This is opposed to the traditional paradigm, stemming from the 1980s, where exact geometric "wire-frame" models of a particular environment are designed by the human supervisors and transferred to a robot's navigation subsystem. It became clear that the unavoidable discrepancies between an internal environment model and the actual world cannot be solved elegantly within this approach. Since realistic environments are not static there is a need for flexible and adaptive navigation methods such as can be found in biological systems. We will therefore postpone the use of absolute Cartesian grids and focus on (1) an approximate navigation modeling system which (2) can ultimately be combined with a definition of behaviors which allow for an escape from local navigation problems.

[1]Appeared in: Proceedings of BNAIC'03, Nijmegen, pp. 275-282

In the following paragraphs the basic mechanisms of navigation in relative simple organisms such as bees, ants and rats will be discussed, incorporating results at the neural and behavioral level.

1.1 Animal navigation

Dead Reckoning, the first mechanism which seems to be involved in animal navigation concerns the continual updating of position and heading by summing successive small ego displacements in the environment. Such an integration of the velocity vector with respect to time will yield an estimate of the position vector. Using this method, the animal is capable to keep track of its current position and directional heading in the environment when there is no perceptual input available to serve as a cue on which it can determine its current position and directional heading. At the neural level, two types of neurons in the rat brain have been implicated in spatial learning and navigation processes: the so called *place cells*[6, 5] of the hippocampus and *head direction cells* of the thalamus and postsubiculum.

Landmark Orientation. Another prevailing mechanism in animal navigation is landmark based. An animal may use the perceptual pattern of familiar, distinctive and easily perceptible points in the environment to determine its current directional heading and position in the environment. The available perceptual input allows for an association with a learned perceptual pattern of a landmark perceived earlier, activating the relative positions of that landmark in the environment to other nearby landmarks which were visited during the learning history. The resulting network of associations is called a cognitive map (Tolman[9]), which in its representation strongly differs from Cartesian grid maps. The use of landmarks orientation is a useful addition to dead reckoning mechanisms, since the detection of a landmark allows for resetting the accumulated position and heading error on the basis of the known and encountered landmark. Wehner and Raber[7] conducted experiments done on desert ants *Cataglyphis bicolor* that suggest that in order to move to a particular location ants try to match their current perceptual input in the form of a visual percept, with stored visual inputs of landmarks. Cartwright and Collet[1] conducted similar experiments on bees, with similar results. At the neurological level of navigation, Muller and Kubie[5] conducted an experiment with rats which provided evidence that *place cells* are also capable of making use of visual landmarks in the environment to establish a place-specific firing pattern.

Other mechanisms for biological navigation exist, such as beacon navigation and gradient-based navigation. These are treated elsewhere[3].

1.2 A simplified model

We will start to explore the concept of cognitive robot navigation based on navigational principles found in relative simple organisms, such as insects, using a Pioneer-II robot. The ultimate end result in our approach will yield a system which is capable of navigation on basis of a landmark-based map of the environment that is being constructed incrementally while the robot is exploring the environment. The navigation model will have to be based on the following components: (a) Multi-modal Landmarks, (b) Dead reckoning and (c) Cognitive maps.

A multi-sensor definition of landmarks. A fundamental characteristic of a landmark is that it is clearly identifiable within the ongoing perceptual stream, and sufficiently unique in the environment to be a basis for navigation. Possibly, a salience feature can be extracted, which indicates landmark candidates in the perceptual stream. It seems likely that such mono-modal perceptual mechanisms do indeed exist (e.g.: a highly satiated color of a flower). However, assuming that landmarks are in some ecological way relevant to the species at hand, it seems likely that multi-sensor input will play a role in landmark determination. Notably, we introduce here the notion of *proximity event*. It is assumed that landmarks are specific samples of multi-sensory input which refer to a location in the environment where 'egocentric space' is invaded by physical objects which are relevant to the interests of the autonomous agent. Proximity events are important because they may indicate imminent collision and/or the presence of Friend/Foe or Food (FFF). Colby and Goldberg[2] describe neurons in the medial intraparietal(MIP) area specialized for responding to stimuli (in our case: potential landmarks) that invade this 'egocentric space'.

Dead reckoning. Dead reckoning is a vulnerable numerical integrator mechanism in many computational substrates. In some animals, optic flow may be the basis of ego-velocity estimation. For the experiments, we could only make use of the Pioneer-II odometry, which appeared to be highly biased by wheel slip. Heading information is especially unreliable, and no vestibular sensor is present to determine actual achieved angular velocity. However, it appeared that the total length of the approximately linear path traveled between two points can be determined with sufficient accuracy.

Cognitive Maps. We think that the cognitive map, as introduced by Tolman[9] on the basis of experiments with rats, are a core component of a working model of biologically plausible navigation. A cognitive map represents the environmental layout. It encodes the metrics(angles, distances) and sense relation(left versus right) between landmarks in the environment. Its representation is vectorial, and different from the 'omniscient' Cartesian-grid based maps. Research at the neural level of navigation suggest that the dead reckoning process is the key to construction of this cognitive topological map in the brain[4].

The first and major goal of this study is to find out whether coupling the visual information to proximity events leads to the development of an unsupervised landmark map which contains the basic elements of the visual and physical environment of the robocup field, using a Kohonen self-organizing map. Second, given a learned cognitive map of the environment including each learned landmark and its distances to other landmarks, the goal is to find out whether this model has been able to capture the global two-dimensional structure of the environment.

1.3 Methods

On the basis of the model requirements described above, a number of design issues have to be solved. Also, the limitations of the used robot platform have to be taken into account. First of all, a suitable visual representation is needed. Tests will be performed in a 1^{st} generation robocup field with white side boards, a blue and a yellow goal, and a green floor. As a first test, the navigation model should be able to detect the goals, the corners and other salient elements in the environment. The proximity event, which is needed to determine that the visual array in front of the robot might be a potential landmark, will be generated on the basis of the ultrasonic sensor readings. The bimodal perceptual landmark vector will consist of a reduced-resolution camera image augmented with appropriately scaled ultrasonic readings.

Image. In the experimental "Robocup" environment, important elements are distinguishable on the basis of color. For this reason, also taking into account the real-time requirements, the use of detailed shape encoding by means of, e.g., Gabor filters, has been postponed. On the basis of such design considerations, a luminance invariant portion of the well-known hue, saturation and brightness (HSB) coding, i.e., Hue-Satiation (HS) space will be used, where the original luminance of the RGB values is normalized to the maximum value. Pilot experiments revealed a desirable insensivity to lighting and weather conditions. An reduced-resolution image of 40x30 (WxH) pixels is used for the Hue and Satiation (HS) feature vector ($N_{dim} = 2400$).

Proximity. The vector of sixteen sonar sensor readings was joined with the image feature vector resulting in a bimodal perceptual input vector, a *perceptual frame* ($N_{dim} = 2416$).

Unsupervised learning for landmark detection. In order to solve the problem of automatic landmark detection and learning, the robot ran a generic collision-avoidance behavior, collecting perceptual frames whenever frontal proximity events where triggered. Considering the amount of visual elements in the environment, it was assumed that for a 5x5 Kohonen self-organizing map (SOM) it should be possible to detect up to 25 'landmarks' in the robocup environment. A 2-D Kohonen SOM was chosen because a planar neural tissue patch is a more plausible solution than its N-dimensional generalization. The advantage of the Kohonen learning scheme is that the nodes in the self-organized map will contain robust average representations corresponding to a set of similar perceptual frames in the learning history. As a result, the landmark detection will be noise tolerant. It should be noted that the topology of the SOM is not coupled to the topology of environment map in a simple way.

In this study it is assumed that salient perceptual patterns of potential landmarks can be collected on the basis of collisions or near-proximity events. Rather than expecting landmark singularities to emerge solely from the training of all visual patterns[10] which are present during all of the ego-movement, the subset of visual patterns at and around a physical collision or near proximity point are deemed essential for the development of a landmark map. Apart from the already mentioned neurophysiological support for the notion of the proximity event, it can be easily understood why multimodality is useful: It provides a strong correlation between separate modalities, concerning a single phenomenon in the environment. This may become more clear if one considers a modality which has a larger bandwidth than proximity-event detection, i.e., tactile sensing: For a looking system, a visual edge often becomes a mechanical edge at the moment tactile contact takes place.

Cognitive map representation. With the extracted perceptual representation of landmarks a cognitive or topological map of the environment will be constructed. As it is the case in animals this cognitive map will represent the relative positions of landmarks in the environment to each other. The distance between landmarks will be collected with a mechanism based on dead-reckoning (odometry). This means that at the end of an exploratory phase, a matrix with distances Δ_{ij} between a landmark i and a landmark j will exist as a representation within the navigation system.

Robot platform and Environment. The robot used for this research experiment is the Pioneer II DX robot. The physical characteristics of the robot are length 44 cm, width 33 cm and height of the body 22 cm. This robot is equipped with an array of 16 sonar transducers that provide range information of objects in the environment from 10 cm to more than five meters. On top of the robot a color CCD camera is mounted with a resolution of approximately 440k pixels. An optical revolution counter on the axles, measuring 19 ticks per traveled millimeter is used for speed sensing and dead-reckoning during linear trajectory segments only. As regards the software platform, the different tasks and behaviors used on the robot in these experiments were implemented on our modular agent-based architecture 'XSAM'. The environment used in conducting the experiments was the *robocup field* in the robotics lab of the AI Institute at Groningen University. This rectangular green field consist of white boarding with a height of approximately 40 cm on each side. The dimension of the field is approximately 4 by 4.60 meters. On each of the two short sides of the field is an colored area that represents a goal. The colors of these goals are yellow and blue. In the robotics lab the lighting is provided by fluorescent ("TL") illumination and a window positioned at the side of the yellow goal. When standing in front of the blue goal there is a darker side with a wall on the right hand and a brighter side with computer terminals on the left hand.

2 Results

After a number of pilot experiments, it was determined that a driving speed of 0.2 m/s, a frame rate of 2 Hz and a proximity threshold of 1 m yielded a useful 'proximity set', i.e., the subset of all perceptual frames consisting of those frames which are collected at proximity events. Experimental trials are of the order of 20-30 min., collecting 2400-3600 frames, of which 500-700 perceptual frames are classified as proximity events. The Kohonen training parameters were: A 5x5 network, using 200 epochs, starting with a radius of 100% and a learning rate of 0.99 and ending with a radius of 0% and a learning rate of 0.01. A steep and non-linear parameter curve for radius and learning rate was used in order to prolongate the relative duration of the final fine-tuning stage. At the end of training, the root-mean square (rms) error between the SOM and the N training vectors can be computed: $\epsilon_{rms} = ((\sum_i^N \sum_j^F (x_{ij} - m_{kj})^2)/NF)^{0.5}$, where $i = [1, N]$ is the sample index, $j = [1, F]$ is the feature-value index and k is the index of the nearest neighbor cell in the SOM, i.e., m_k, to the input sample x_i. The rms error value for four networks, trained on four days with different weather and light conditions varied from 25-41 and there were no appreciable differences between same-day and different-day cross comparisons, using the HS images. As expected, rms errors were larger for an RGB image representation, using the same scaling (0-255). Figure 1 shows the resulting SOM, or rather: Kohonen Landmark Map (KLM), after training on HS images. The images on the right represent the sets of raw feature frames which share the same nearest neighbor in the KLM.

In the visual Kohonen Landmark Map (Figure 1) patterns of important landmarks in the environment such as the blue and yellow robocup goal, which are easily perceptible and unique, clearly emerge (in grey-scale print, the blue goal looks dark). The raw feature frames which have one of these landmark representations(blue and yellow goal) as the same nearest neighbor in the KLM show a clear homogeneity. Landmarks representing points in space which are not easily perceptible or unique on the visual pattern show more heterogeneity in the set of feature frames with the same nearest neighbor. For example corners and the points on each side of the field reveal great similarity in the visual domain, whereas the effective use of the sonar information may help to disambiguate[3].

At this point, it is an important question whether the amount of information collected by the system thus far is sufficient to construct an approximate map of the 2-D environment. Given a set of N points in the 2-D plane, an NxN distance matrix can be constructed. It can be easily shown that it is possible to estimate the position of a point i in this plane on the basis of a linear transform on its distances to the other $N-1$ points,

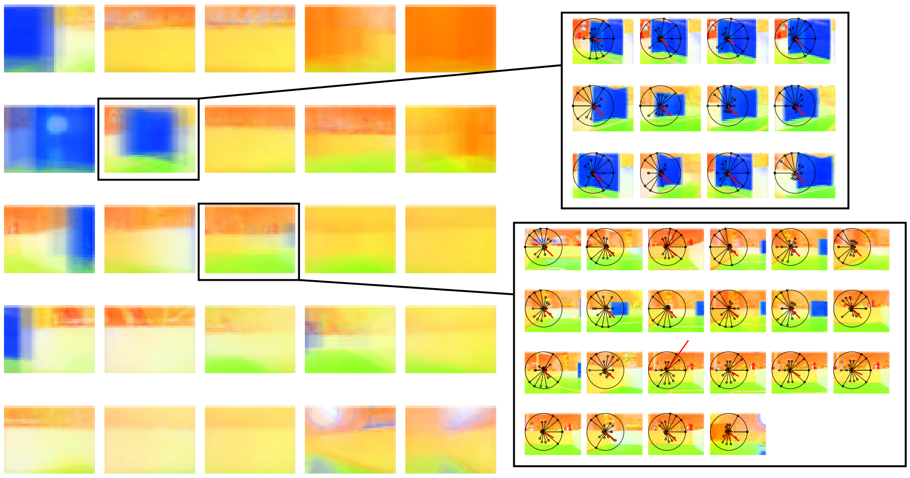

Figure 1: Left : KLM after training on HS images; Right : sets of raw feature frames belonging to node (2,2) and node (3,3) respectively as the nearest neighbors in the KLM. The main image of the blue goal is at row 2, column 1. The main image of the yellow goal is at row 1, column 5 (1-based indexing is used here). The black circle and arrows in each cell represent the average sonar data as $1/r$ (long='near') belonging to the landmark image. It can be seen that the image data dominate the sonar patterns.

by solving a set of linear equations and using known coordinates as the desired output:

$$\hat{x_i} = \beta_0 + \beta_1 \Delta_{1i} + \beta_2 \Delta_{2i} + ... + \beta_N \Delta_{Ni} \quad (1)$$

$$\hat{y_i} = \gamma_0 + \gamma_1 \Delta_{1i} + \gamma_2 \Delta_{2i} + ... + \gamma_N \Delta_{Ni} \quad (2)$$

where β and γ refer to the linear coefficients for x and y, respectively.

In other words, according to this model, the distance matrix Δ contains information on the dimensionality and geometry of the space containing the N points. By determining the Eigenvectors of Δ, an unsupervised estimate of the 2-D structure can be obtained theoretically, on the basis of the two axes with the largest Eigenvalues, using principal-components analysis (PCA). The resulting representation will need to be normalized on the basis of an affine transform (4 parameters).

To illustrate this principle, an artificial distance matrix was constructed which represented the distances between 25 points uniformly distributed in a 5 by 5 square. One percent of the distances was not realized and replaced by a random value within the same range as the distance values. On all distances, uniform noise with a maximum of 0.3 was added. Figure 2 shows the grid points and the scaled estimates which are based on the PCA method. As can be seen, the reconstruction is not without error, especially when considering points at the perimeter of the field, but the general layout is reconstructed. Noise on the distances in the matrix will increase the error in a reconstruction of the map somewhat. Missing distances have a stronger bias effect. This result on artificial data illustrates, that a landmark-exploring agent can construct an approximate vectorial map of a limited environment, in principle.

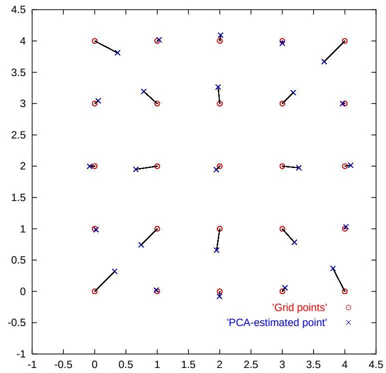

Figure 2: PCA results on artificial created distance matrix between 25 points, with 1% missing links and a white noise of max. 0.3 on the distances, in this scale. Circles mark the original 2-D grid points, the x crosses represent the estimated grid points according to the two largest Eigenvectors of the distance matrix.

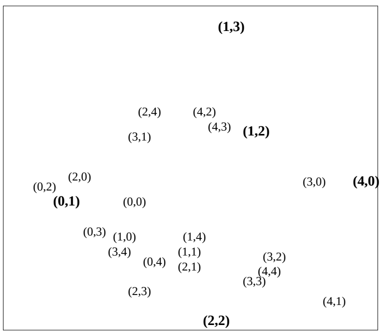

Figure 3: Topological landmark representation based on PCA of the distance matrix. Without human intervention, the system has correctly place the blue(0,1) and yellow(4,0) goals in opposing positions. The landmark views on the room's window, i.e. nodes [(1,2),(1,3)] at the top, indicate that the y-axis should be flipped in order to obtain the real-world configuration.

In order to test this model on real data, the robot ran through the environment, using a given 5x5 KLM and noting both the identities of each landmark image and its dead-reckoned distance to the preceding landmark. From the resulting data a 25x25 distance matrix was filled. Distances for landmark transitions which did not occur during the test run were replaced random noise with mean and standard deviation of the known distance distribution. Figure 3 show the topological landmark representation of the environment which resulted after PCA on the distance matrix. The labels in the topological map correspond to the labels in the perceptual landmark of the KLM used. The map represents the two main axis of the robocup field: Blue opposite to the yellow goal and 'terminals' side opposite to the 'dark wall' side.

3 Conclusion

We have proposed a biologically inspired landmark-exploration model, using multimodality and unsupervised learning. Furthermore, it was illustrated that a landmark-distance matrix - such as may be determined on the basis of dead reckoning - contains the necessary information for the development of a vectorial 2-D map. By using principal-component analysis, the two major eigenvectors were illustrated to refer to the major Cartesian axes of an environment, both for artificial grid data and for real exploratory navigation on a robocup field. However, it should be noted that more research is needed, since the model is limited in a number of ways. The fixed dimensionality of a Kohonen map does not seem to be biologically plausible. Furthermore, a calibration of the PCA-based cognitive map is needed, requiring an affine transform (4 parameters) in order to scale and rotate the internal map appropriately relative to motion control. In principle, also this last step may be unsupervised, if tuning behavior is added to the navigation system. **Future research** will be directed at an extension of the model, using behavioral strategies. For example, at a proximity event and a detected landmark, looking left and right will allow for a more unique coding of a landmark, coupling it more strictly to the geographical location in the real world, but also providing for relative angular information which couples the landmark to its neighbors during the common exploration behavior.

References

[1] B.A. Cartwright and T.S. Collett. Landmark learning in bees. *Journal of Comparative Physiology*, 151:521–543, 1983.

[2] Colby C.L. and M.E. Goldberg. Space and attention in parietal cortex. *Annual Review of Neuroscience*, 22:319–349, 1999.

[3] R. Fehrmann. Cognitive navigation modeling in robots. Technical Report [MSc Thesis], AI Inst./Rijksuniversiteit Groningen, 2002.

[4] B.L. McNaughton, C.A. Barnes, J.L. Gerrard, K. Gothard, M.W. Jung, J.J. Knierim, H. Kudrimoti, Y Qin, W.E. Skaggs, M. Suster, and K.L. Weaver. Deciphering the hippocampal polyglot: The hippocampus as path integration system. *Journal of Experimental Biology*, 199:173–185, 1996.

[5] R.U. Muller and J.L. Kubie. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7:1935–1950, 1987.

[6] J. O'Keefe and J. Dostrovsky. The hippocamplus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, 34:171–175, 1971.

[7] Wehner R. and F. Raber. Visual spatial memory in desert ants. *Cataglyphis bicolor*, Experientia 35, 1979.

[8] S. Thrun, W. Burgard, and D. Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping. In *Proc. IEEE Conf. on Robotics and Automation ICRA*, San Francisco, CA, 2000. IEEE.

[9] E.C. Tolman. Cognitive maps in rats and man. *Psychological Review*, 55:189–208, 1948.

[10] Georg Von Wichert. Selforganizing visual perception for mobile robot navigation. In *1st EUROMICRO Workshop on Advanced Mobile Robots*, 1996.