"NAVBOT" - Autonomous Robotic Agent with Neural Learning of Autonomous Mapping and Navigation Strategies

Pedro Kulzer

Resumo - **Neste trabalho, foi realizada uma implementação dum agente robótico autónomo minimalista com capacidades de navegação avançadas, as quais dependem em arquitecturas especializadas com base em redes neuronais artificiais. Pretende-se mostrar que uma plataforma robótica comparativamente simples e dotada de componentes baratos é capaz de realizar tarefas complexas de navegação, tais como *dead-reckoning* (intuição espacial), circunscrição de pistas visuais e correspondente discriminação / reconhecimento, bem como a construção eficiente de mapas e seu uso para navegações futuras com a tomada de atalhos e voltas. São usados *shaft-encoders* para a implementação de uma bússola interna grosseira, bem como sensores de infravermelhos para a detecção e circunscrição de objectos. As redes neuronais utilizadas incorporam informação espacial bem como temporal. Antes de realizar a implementação final, foi realizada uma pesquisa bibliográfica inicial, a qual permitiu saber o que existe de actual neste campo da navegação em agentes autónomos.**

Abstract - **In this work, an implementation of a minimalist autonomous robotic agent with advanced navigation abilities was done. These abilities rely on specialised architectures based on artificial neural networks. It is intended to show that a comparatively simple robot platform with cheap components, is able to perform complex navigation tasks such as dead-reckoning (navigation through spatial intuition), landmark circumvention and corresponding discrimination / recognition, and efficient map construction and use for future navigation with shortcuts / detours. Shaft-encoders are used for the implementation of a coarse internal compass, whereas infra-red sensors are used for landmark detection and circumvention. The involved neural networks incorporate space as well as time information. Before doing the final implementation, an initial bibliographic research allowed to know the state-of-the-art in this field of navigation in autonomous agents.**

I. INTRODUCTION

This paper results from a Master's thesis work described in [7]. Landmark/place perception, recognition, storage, and subsequent navigation with a map, are a very interesting field of research that seems to be still traversing its birth. Many attempts have been made until today in the field of self-localisation and mapping, ranging from more theoretical studies which describe the efficient and powerful storage of spatial information in the rat's hippocampal brain structure as a sequence of places identified by sensory combinations with the aid a "clocking" brain wave called the *theta rhythm* [14] (see illustration in Fig. 1) [17][18], passing to simulated studies [1][2][16][24], reaching real implementations and experiments where an artificial agent tries to store and recognise local spatial distance and turn angle information of a landmark by its simple circumvention [3][4][5] (see illustration in Fig. 2) [8][19].

All these studies mix the concepts of landmarks and places, without really specifying what distinguishes them and which one should be used. Few of these studies incorporate the notion of dynamic place sequencing issues [14][17]. The remaining ones, just talk about static models, where recognition is done by means of independent environment snapshots.

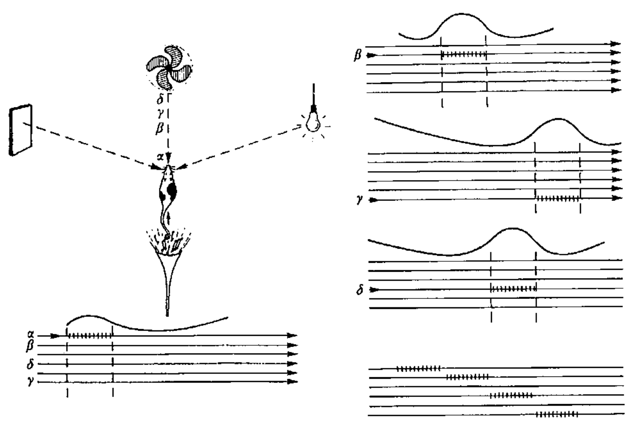

Fig. 1 - Model of *Hippocampal* place sequence storage. (adapted from [14])

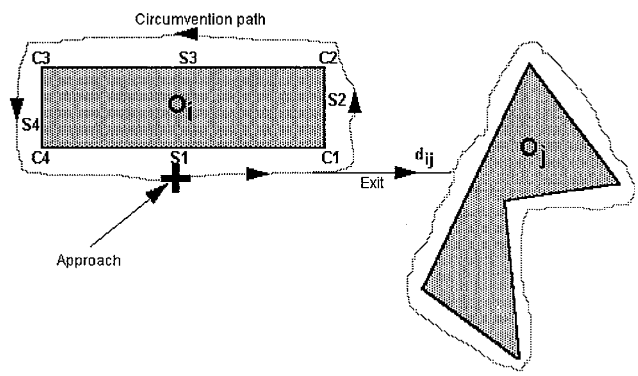

Fig. 2 - Landmark recognition by global circumvention. (adapted from [19])

The main question that this work tries to answer, is to see if using only turn angle information and sequencing is sufficient for robust landmark shape recognition. Also, it tries to establish a relationship between a landmark and its corresponding places around it.

The major guidelines during this work, were:

- Instead of strict AI with *if-then-else* structures, neural mechanisms will be used extensively.
- Always try to satisfy the maximum amount of experimental data with the developed models.
- Try to interconnect and correlate different experimental data and theories into one global theory of recognition and mapping mechanisms.
- Prefer models and theories that may be implemented in parallel and distributed *computer hardware.*
- Prefer ***on-line*** learning mechanisms, instead of such where complete data must be present at once.
- Always try to make every value relative to another, so cumulative global errors are reduced.
- Keep hardware cost as low as possible, using cheap sensors, processors, robotic platform, etc.

II. THE HIPPOCAMPUS AND RELATED AREAS

Processes of mapping and navigation are thought to be highly related to a brain area called the *Hippocampus.* Neuroscientists showed that there are neurons responsible for signalising upon the detection of the animal's position in a certain environment [9][10][11][12][13][14]. There are neurons, called *place neurons*, whose electrical activity is directly related to the animal's movement in particular locations of the environment. Conversely, there are the so-called *place fields* in the environment that correspond to the firing places of particular place neurons, in that particular environment.

There are *theta cells* in the hippocampus that exhibit a firing pattern similar to the *Theta-Rhythm* of the hippocampal EEG. This rhythm is best characterized by a sinusoidal wave with a frequency that can reach from 4 to 12Hz [14] and appears to be roughly proportional to the animals' motion speed. It appears only when the rat is moving around, but not when it is doing stationary things like grooming, eating and sitting still [22]. The destruction of the brain area responsible for the generation of this rhythm 'pacemaker', results in a loss of the spatial problem solving capabilities of the rat [23].

On the hippocampal map model shown in Fig. 1, it is postulated that the theta rhythm is part of a mechanism for shifting the focus of excitation (learning) from one set of place representations to another within the map, on the basis of the animal's actual or intended movements. They also found this theta rhythm seems to provide the distance magnitude information for "internal navigation". The direction information could either be given by external angles to landmarks (exteroceptive) or either by an internal compass or calculations based on the turns made by the animal (interoceptive). This way, different neuron groups get sensitised [15] to different instants in a spatial path, forming a *sparse memory representation* of that path (each memory element is the sensory combination at each instant). Thus, the path from α to δ will be systematically stored as the sensory representations of the corresponding sequential places.

After the exploration period, the animal then generates a low-frequency theta-rhythm scanning across the neuron groups until the representation of place α is activated. This activation is then the "go" signal for the animal to move towards β and also served to synchronise the theta rhythm within the map.

Greater movements across greater distances are accompanied by higher theta frequencies, so that the phase relationship between places in the neuron group representation are still the same.

The *parietal cortex is* probably responsible for the storage of metrical information about relations between places, by receiving metrical data from the *motor cortex.*

Important notions of *spatio-temporal self-organising feature maps* are enhanced in [6] for pattern sequence detection and recognition. This recognition is triggered by the growth of an increasing wavefront of accumulative activity. While learned sequences produce an accumulation of coherent neural activity, other unknown sequences produce a lower and non-coherent activity progress. This basic idea will be used in the present work.

III. THE ROBOTIC PLATFORM

The robotic platform is based on a previously assembled robotic platform called "THOMAS" [20]. THOMAS was specifically designed for the purpose of landmark circumvention. It has been improved in [7] and had the following main hardware features (see Fig. 3):

7 IR sensors - 2 to the left and right, and 3 to the front.
2 Motors - 2 motors at the front and 1 caster wheel.
2 Shaft-encoders - 1 shaft-encoder per wheel.
1 CPU board - 68HC11 microcontroller plus software.

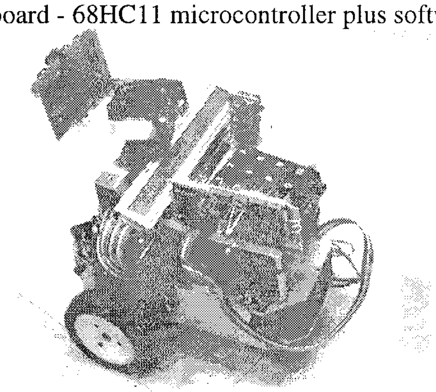

Fig. 3 - The real robotic platform "NAVBOT". (adapted from [7])

In addition, the robot is able to keep an internal compass, and compute displacement vectors between locations. All this is done by very simple vector computations with the help of the incremental shaft-encoder distance counts for each wheel.

IV. The model

As in [3][4][5][19], the present model and implementation is based on landmark circumvention. This is similar to what a blind person does when touching objects. This model has the enhancement of being able to recognise any shape, and not only polygonal ones as in previous work. In earlier implementations, when the robot missed a corner, then the whole recognition process might be wrong. Even if a previously undetected corner now appeared, the same problem might arise. This is because the relative distances between recognised corners then appears to be very different from the previously learned ones during exploration.

In the present robot, IR sensors act as short distance proximity sensors for *wall-following*. Shaft-encoders allow the robot to measure direction changes, specially for general turn angle measurement. Several stages of processing occur when the robot is recognising a landmark. *Finding* occurs when the robot perceives an object, where it will dock parallel to it. After this, the robot will start the *circumvention* around the object, by means of *wall-following*. We chose to perform wall-following as the recognition data gathering mechanism, since this was the most the robot could perform with its limited sensing capabilities.

A. Data storage

While wall-following the landmark, the robot must somehow gather data about the landmark as seen by the robot. This data is definitely related to the landmark's shape (see Fig. 6). If only the corners and distances between them are taken as classification data, then only polygonal landmarks could be recognised, the matching algorithm would be brittle, and the corner extraction itself would be brittle. To avoid these problems, a much different approach must be taken.

The need for a neural mechanism whose matching capabilities are invariant to the circumvention starting point, leads to a somewhat changed neural net. A very simple polygonal example of what is going to be done follows, just to keep explanations simple (Fig. 4). Everything can be extrapolated to any landmark shape. Only local relational values are taken, namely the turn angles between each pair of distance segments. These segments are kept uniform in size, which means that the robot samples at equidistant points. The corresponding stored networks are shown beside each landmark. Here, each synaptic circuitry stores the angular turn between corresponding places. Each "turn" is regarded as the angle formed by the last two displacement segments performed by the robot.

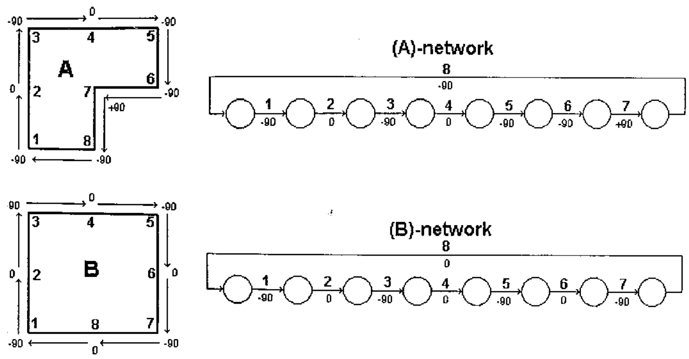

Fig. 4 - When the robot circumvents a landmark for storage, it generates a neural network with the successive turn angles.

Note the similarity of the successive place neuron sensitisation along the circumvention, with the theta rhythm sensitisation action in the hippocampal neuron groups discussed earlier. The shaft-encoders, besides giving information about the turning angle, also give simple displacement information to move to successive neurons in the network.

B. Recognition

The major problem with this new model is related with the sampling positions. The robot learned a network with weights that represent the turn angles. If the robot circumvents the landmark starting at an intermediate position, all angles will be "out-of-phase" (see Fig. 5).

To solve the problem, an *averaging motor-cortex* is implemented, as a preprocessor for the neural network. This mechanism that takes an *averaging window of motor activity* to compute the segments that are going to be fed to the recognition network. Each motor-cortex neuron has its maximum response at a certain turn angle. The overall resulting system is depicted in Fig. 7.

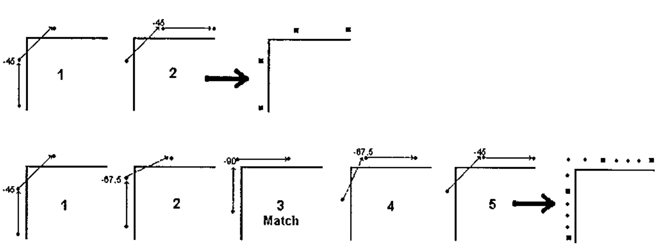

Fig. 5 - Successive overlapped segment sampling, where the stored angles periodically match.

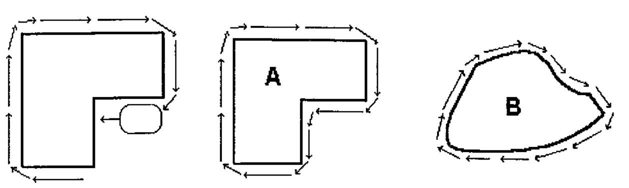

Fig. 6 - Wall-following around any shape. This work tries to show that even a sloppy circumvention is enough for the neural mechanism to operate.

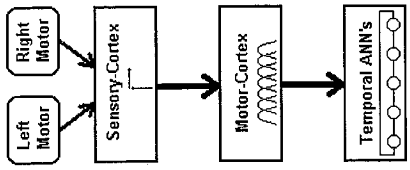

Fig. 7 - Motor-cortex tuning-curve outputs and resulting system used for circumvention data storage and recognition.

Fig. 10 and Fig. 8 show the recognition mechanism used, which is similar to the activity wavefront propagation in [6], The difference is that the activity starts at 100% and decreases with successive turn angle mismatches. This example shows one wavefront keeping 100% activity, while the other suffers mismatches and dies out, for landmarks A and B, respectively (see Fig. 4). Note that larger mismatches cause the multiplicative factor to be smaller. All wavefronts start at 100% from each place neuron and suffer successive mismatches and matches which, in turn, propagate more or less activity. The surviving wavefront is called the *principal wavefront* (expected correct matching sequence), while other less active and still persistent ones are called *phantom wavefronts*. Because in this example the robot is trying to recognise landmark A (see Fig. 4), all wavefronts in landmark B die out sooner or later. Note that they really start to die out as soon as the robot sees the corner sequence -90° +90° ... which differs from landmark A.

Note that the robot may start at any point, where different points produce the survival of different principal wavefronts. This means that the robot is even able to know where it currently is on the landmark, while circumventing it. More precisely, the current principal wavefront's index in the network corresponds to a place neuron firing maximally which, in turn, directly corresponds to a precise place around the landmark. This place changes as the firing neuron also changes as the wavefront propagates on.

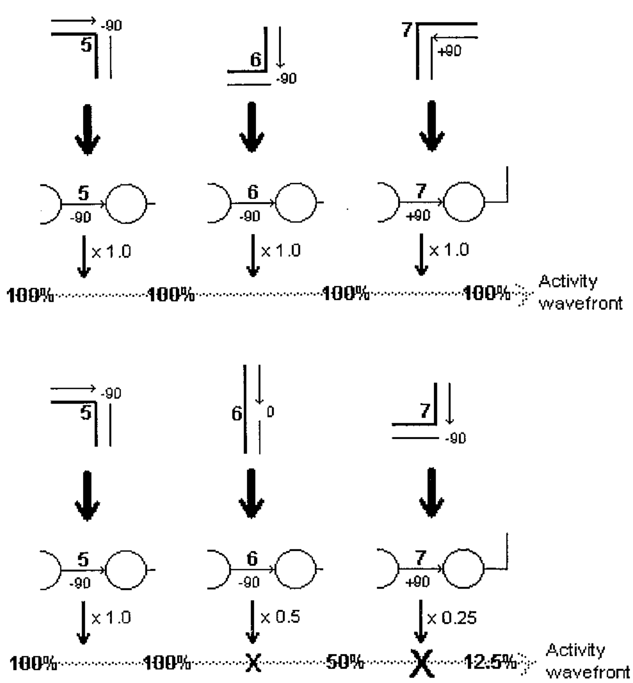

Fig. 10 - Selective activity wavefront propagation.

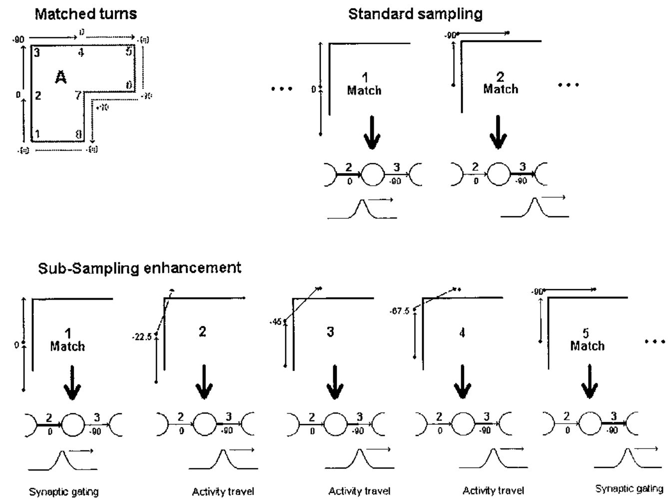

Fig. 8 - Details of the theta rhythm aided matching.

C. *Real tests*

Some real tests were performed on three landmarks (see Fig. 9), to see the discrimination, recognition, and robustness of the designed neural mechanisms.

Tab. 1 lists some of the many circumvention trials performed. There were some major problems with the stopping procedure which consisted in detecting an overall displacement vector below a low threshold value. Neuron counts varied within a short range. Fig. 11 shows some graphical reconstructions of the corresponding landmark shapes (asterisks were the starting points). Although their global shape is distorted, the network works well on the relative turn angles. To simplify, we show only the details for the first landmark.

Trial	Turn angle sequence	N
Run1	67.5 - 0 - 0 - 90 - 22.5 - 0 - 0 - 45 - 22.5 - 0 - 67.5 - (-22.5) - 0 - 22.5	14
Run2	67.5 - 22.5 - 0 - 45 - 45 - 0 - 0 - 67.5 - 22.5 - 0 - 22.5 - 67.5 - 0	13
Run6	67.5 - 22.5 - 0 - 45 - 45 - 0 - 0 - 45 - (-22.5) - (-22.5) - 45 - (-22.5) - 0 - 0 - 45	15
Run12	45 - 0 -0 - 22.5 - 67.5 - 0 - 0 - 45 - 0 - 0 - 22.5 - 67.5 - 0 - 22.5	14

Tab. 1 - Stored angle sequences for the test runs.

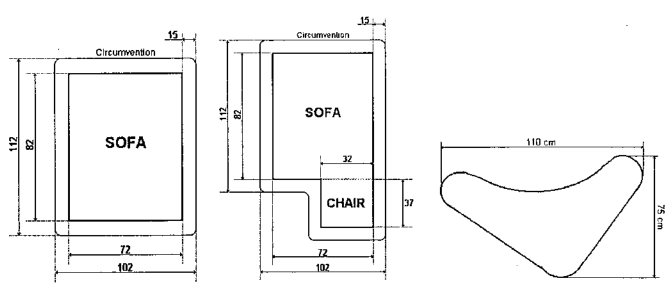

Fig. 9 - Three real test landmarks.

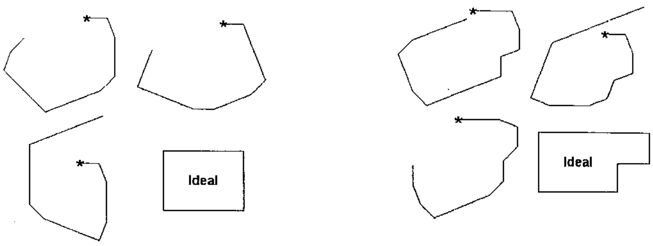

Fig. 11 - Some robot-data reconstructed landmark shapes for the first two landmarks, along with the ideal reconstruction.

Now, to test preliminary recognition power of the networks, some other circumventions were made, where now six sequences were recorded (Recognition 1 through 6) and fed into the neural networks. Network activity ratios between networks were calculated and compared (see Tab. 2). The smallest activity was 0.49 and the largest was 0.90. Note that these values can be looked as being certainty values, since they vary from 1.0 (100% sure) and 0.0 (0% sure).

	Rec 1	Rec2	Rec3	Rec4	Rec5	Rec6
Run1	0.75	0.76	0.80	0.87	0.90	0.82
Run2	0.62	0.56	0.59	0.62	0.49	0.70
Run6	0.52	0.74	0.65	0.56	0.69	0.51
Run12	0.90	0.82	0.90	0.80	0.82	0.80

Tab. 2 - Experimental recognition results.

The propagation equation is shown in equation (1) where α=0.6, and where *Out* is the propagated wavefront on a neuron with weights *link* and input *input*. Note that the tuning-curves used are the positive part of cosines. In other words, if the difference between input turning angle and current neuron synaptic link value is null, then the match is 100%. If not, then the match value will be smaller with larger differences.

$$Match = \max\left[0, \cos\left(\alpha \cdot (Link - Input)\right)\right] \qquad (1)$$
$$Out = Match \cdot Wavefront$$

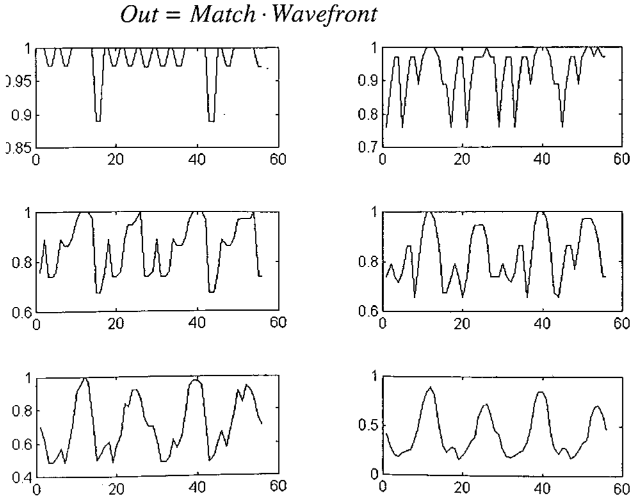

Fig. 12 - Each individual wavefront's activity progress from the 3rd sample to the 50th sample point (whole circumvention).

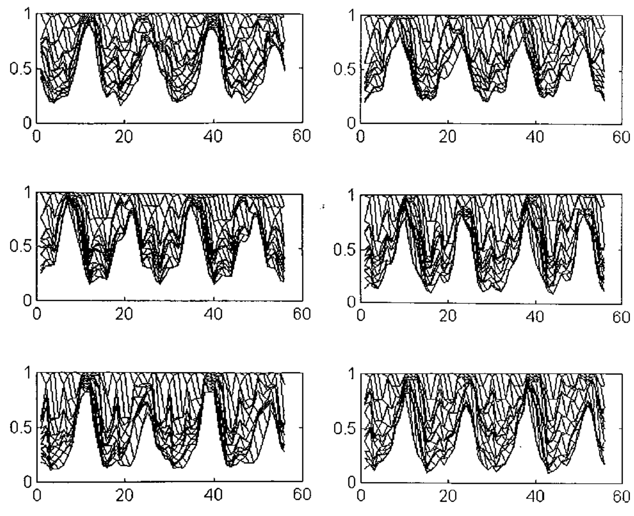

Fig. 13- Experimental results for 6 test runs on the first landmark.

Fig. 12 shows the activity progress graphs, only for the first rectangular landmark. Note that, since the landmark was almost a square, four principal wavefronts will arise. The interesting thing is that there are two with more activity than the other two. This means that the robot was actually able to discriminate between the two pairs of symmetric positions in the rectangular landmark. In other words, it was able to recognise it as being a rectangle.

Fig. 13 shows final progressions, where some show confusion of the robot regarding the real rectangular or square shape of the landmark.

Fig. 14 and Fig. 15 show the details obtained by the circumventions around the second non-symmetric landmark. One can see that only one principal wavefront more or less consistently appears to form. This clearly means that the robot is able to identify more or less precisely where it currently is on the landmark. In contrast to the first landmark, the robot recognises the places it passes, unequivocally. Note that sometimes there are still another three phantom wavefronts noticeable, which may correspond to the other three strong changes in angular turn.

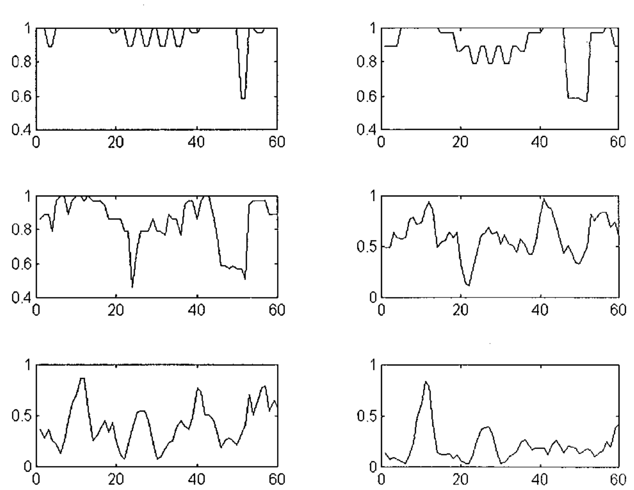

Fig. 14 - Results for the second landmark. (Same as for the first in Fig. 12)

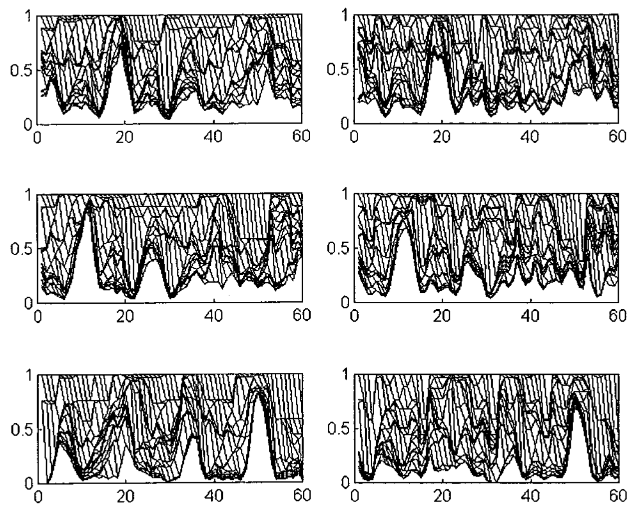

Fig. 15 - Results for the second landmark. (Same as for first in Fig. 13)

The third landmark has curves as well as segments (see Fig. 9). The recognition method will map curves just as it maps corners and segments, with no special distinction at all. Fig. 16 shows that as in the second landmark, the robot is also capable of unequivocally identify the places around this third landmark. Note that now only three major wavefronts appear. It seems that the robots really detects combinations of side + corner with preference, or maybe strong angular turns.

Tab. 3 shows the activity ratios for different stored links and recognition sequences. For example, L_{A12}/L_{B5} - S1 means that the stored links 12 from landmark A (rectangular) and 5 from landmark B (second real landmark) were used, and tested with the recognition Sequence 1. The leftmost column indicates which landmark generated this test sequence, either A or B. For

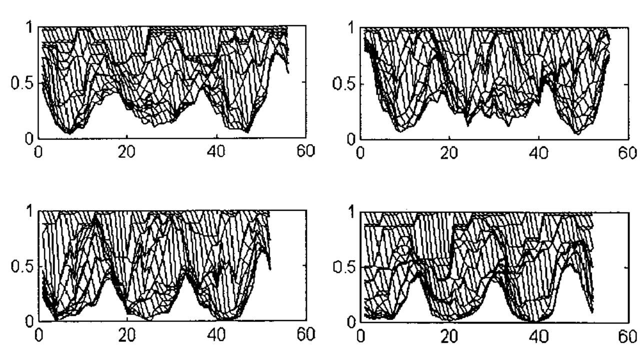

Fig. 16 - Results for the third landmark.

example, a ratio of 1.51 in the first table cell, indicates that the activity from net A was 1.51 times larger than the one from net B (recognition successful).

There are some bad recognitions (bold), essentially due to the already mentioned stopping point problems (see also), giving an error rate of 18.8% and 12.5% for the top and bottom combinations, respectively. This seems to say that the third landmark is easier to discriminate from the others.

Fig. 20 shows some of the recognition ratio progresses as the robot circumvents a landmark and feeds the neural network. Note that the ration just starts increasing after a while. This is due to the similarity between both landmarks until a certain point which depends on the circumvention start position. Observe also that the ratio increases up until a certain point where it may decrease very much. By inspecting the corresponding stored turns and recognition sequences, one concluded without a doubt, that this is due to the length mismatch between the two. One partial solution to minimise this problem is to accept a recognition as soon as the ration becomes "large enough". This can cause errors where the ratio starts going down and then recovers.

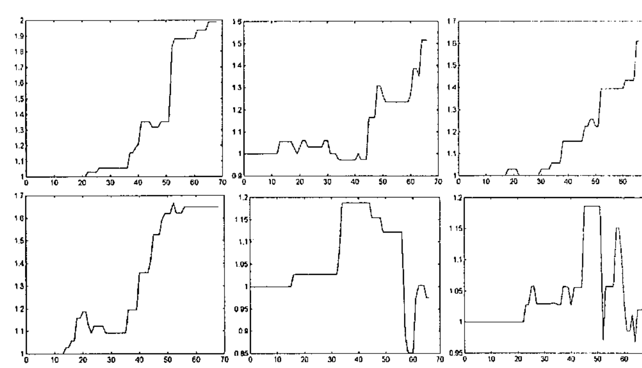

Fig. 17 - Recognition ratio progresses as the robot circumvents a landmark.

Experience showed the evidence of head and tail mismatch of the circular neural network. After the first circumvention around the rectangular landmark, the network is still working well (4 active wavefronts). However, when the robot tries the second circumvention, the wavefronts die out rather quickly. This is due to the frontier effects produced by the mismatch of the end and beginning of the stored turn angle sequence (lateral wavefronts do not catch this deviation any more). It is easy to observe that activity wavefronts die out and reappear in deviated positions (some additive propagation modifications were made just to allow this experiment), either to the right or to the left, depending on whether the stored sequence was too long or too short.

	L_{A1}/L_{B1} S1	L_{A1}/L_{B1} S2	L_{A2}/L_{B2} S1	L_{A2}/L_{B2} S2	L_{A6}/L_{B6} S1	L_{A6}/L_{B6} S2	L_{A12}/L_{B5} S1	L_{A12}/L_{B5} S2
Landmark A	1.51	1.65	**0.98**	1.13	**0.90**	1.23	1.61	1.99
Landmark B	1.34	1.37	1.40	1.32	1.22	**0.90**	1.35	1.20
	L_{A1}/L_{C1} S1	L_{A1}/L_{C1} S2	L_{B1}/L_{C1} S1	L_{B2}/L_{C1} S2	L_{A2}/L_{C2} S1	L_{A2}/L_{C2} S2	L_{B2}/L_{C2} S1	L_{B2}/L_{C2} S2
Landmark A / B	2.78	2.5	1.37	2.18	**0.97**	**0.86**	1.51	1.42
Landmark C	1.05	1.33	1.3	1.25	1.4	1.77	1.3	1.39

Tab. 3 - These results show recognition ratios for the first two landmarks (top) and combinations of the three landmarks (bottom).

V. MAP BUILDING

To finalise, some theories are described. In these theories, a 2D type of representation was used, no spatial nor landmark count limits were imposed, a highly structured environment with, stationary landmarks was assumed, landmarks were assessable, no global North was used, and no other obstacles were present.

Basically, the robot would build global maps by storing pairs of landmarks with corresponding "bi-directional" polar vector links $\left\{\vec{d}_i;\alpha_i\right\}$ in-between, characterised by a displacement and an angle relative to a local landmark reference (see Fig. 18).

Map navigation would be achieved by local neural activity resonance mechanisms among groups of hippocampal place neurons (see Fig. 19) which enable the robot to spread activity from the "conscientiously" selected desired goal landmark to the current landmark which, in turn, enables dynamic and opportunistic path following. A neural activity "resonance path" is formed (see Fig. 20), where closed loops are inherently eliminated. Similar speculations about these mechanisms exist in related models [8].

This model, along with some other details of operation, really allows all the operations observed in the biological case, and are compared to the animal's capabilities:

- **Fast build-up** - the robot would be able of building up the map as fast as it explores the environment.
- **Shortcuts & Detours** - the robot would be able to efficiently select shortcuts and detours. Shortcuts are readily computed by polar vector summation and rotation [21].
- **Goal concentration** - the robot would reach a goal by "thinking" about it (activating the goal neurons).
- **Expectancy** - the robot would have a "feeling" measure of being close to certain nearby landmarks.
- **Failure & Reacquisition** - the robot would be able to reacquire its desired path, after a landmark removal or after having lost itself.

Fig. 19 - Illustration of the resonance mechanism that may happen among groups of hippocampal place neurons which represent places.

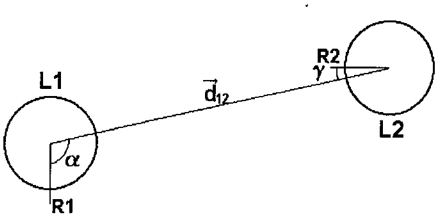

Fig. 18 - Bi-directional link between a pair of landmarks.

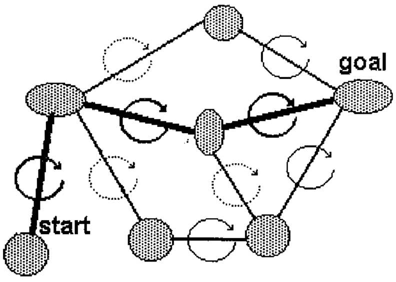

Fig. 20 - Local Resonance forms the best possible trajectory to the desired goal.

VI. DISCUSSION

The main features of this recognition model are the following: no classical Artificial Intelligence *if-then-else* structure whatsoever. The landmark can be of any shape, not only polygonal. Neurons are uniformly created while the robot circumvents the landmark. The synaptic circuitry of each created neuron contains angular information. The need for space-warping is eliminated through the use of uniform spatial sampling. The matching process is straightforward and needs only to select the network with the largest activity. Turn angle deviations are well tolerated (except for head-tail mismatches, of course). Structure and propagation mechanisms are very similar to a *Hidden Markov Model* in respect to the recognition.

Comparisons at the level of the recognition network may be made with the *Hippocampus*:

- **Place neurons** - like in the hippocampus, the synaptic circuitry of these similar place neurons acquire place information, in this case only turning angle data around a landmark, instead of more complex sensory data [14]. The position where the robot currently is, elicits the peak of the principal wavefront at the corresponding place neuron.
- **Theta rhythm potentiation** - like the biological theta rhythm, which seems to successively sensitise the synaptic circuitry of place neurons in the rat's hippocampus here synchronisation is also made by distance sampling.
- **Theta rhythm phase** - phase progress along the place neurons is relatively proportional with respect to the robot's distance travel.
- **Landmark arrangement** - similarly to the spatial arrangement of landmarks in the *hippocampus*, these networks change responses if the landmark changes its turning angles.
- **Movement kinematics** - these networks rely only on turning angle information between positions and

work only robot really moves (wave-front progression through the place neurons).

The overall architecture can be viewed as an *"analog shift-register"* which shifts voltages (wavefronts) according to the gating values (matching).

From all the above considerations and experiments, we conclude that, instead of recognising a landmark as is, we are really recognising sequences of places around that landmark. The robot is therefore able to recognise the current place (assuming a non-symmetric landmark shape) just by the past turn angle sequence. The problem now lies in whether the robot is able to discriminate correctly between these sequences on different landmarks.

VII. CONCLUSIONS AND FUTURE WORK

This landmark recognition mechanism could be viewed as being a microscopic version of the biological mapping mechanism in animals. In other words, instead of storing and relating place fields, this circular network takes a smaller place sequence around a landmark. If the sampling positions are places, then we have a mechanism that closely implements a similar model as in [14]. This kind of network could also be used in macroscopic mapping, where the robot could recognise sequences of landmarks, when lost and trying to reacquire its current position. Still related to animal mapping, one could look at the travelling wavefronts as being *context* information. In other words, when a wavefront reaches a certain neuron inside the network, this means that this neuron receives an indication that it should fire next (expected result in that context).

Future work could eventually focus on achieving better results and turn this architecture more biological, one could try using *sensory constellations*, i.e., more sensors that combine into outputs that somehow reflect specific sensory cue configurations is space. Something similar to [14] could be attempted, specially the sensory cortex output lines. Hopefully, places around the landmarks would be more discriminative and lead to easier to discriminate sensory constellation sequences.

I also think that, in general, the landmark recognition issue must be reconsidered, to see what exactly should be done at the level of the robot's self-localisation.

REFERENCES

[1] N. Burgess, J. O'Keefe, M. Recce, "Using hippocampal 'place cells' for navigation, exploiting phase coding", *Advances in Neural Information Processing Systems*, 5, 929-936, 1993.

[2] N. Burgess, J. O'Keefe, M. Recce, "A model of hippocampal function", *Neural Networks*, 7, 1065-1081, 1994.

[3] S. Caselli, K. Doty, R. Harrison, F. Zanichelli, "Mobile robot navigation in enclosed large-scale space", Machine Intelligence Laboratory, University of Florida.

[4] K. Doty, S. Caselli, R. Harrison, F. Zanichelli, "Landmark map construction and navigation in enclosed environments", Machine Intelligence Laboratory, University of Florida.

[5] K. Doty, S. Seed, "Autonomous agent map construction in unknown enclosed environments", *MLC-COLT '94 Robot Learning Workshop*, 1994.

[6] N. Euliano, J. Príncipe, "Spatio-temporal self-organizing feature maps", to be published in *Proceedings of the ICNN 96' Conference*, 1900-1905, 1995.

[7] P. Kulzer, "NAVBOT - Autonomous robotic agent with neural network learning of autonomous mapping and navigation strategies", Master's Thesis, University of Aveiro, Portugal.

[8] M. Mataric, "Integration of representation into goal directed behavior", *IEEE Transactions on Robotics and Automation*, 8, 304-312, 1992.

[9] B. McNaughton, "Neuronal mechanisms for spatial computation and information storage", *Neural Connections, Mental Computation*, 285-350, 1988.

[10] R. Muller, J. Kubie, "The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells", *Journal of Neuroscience*, 7, 1951-1968, 1987.

[11] R. Muller, J. Kubie, J. Ranck, "Spatial firing patterns of hippocampal complex-spike cells in a fixed environment", *Journal of Neuroscience*, 7, 1935-1950, 1987.

[12] J. O'Keefe, "Computations the hippocampus might perform", *Neural Connections, Mental Computation*, 225-284, 1989.

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

[14] J. O' Keefe, L. Nadel, *The Hippocampus as a Cognitive Map*. Clarendon Press, 1978.

[15] C. Pavlides, Y. Greenstein, M. Grudman, J. Winson, "Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of θ-rhythm", *Brain Research*, 439, 383-387, 1988.

[16] M. Penna, J. Wu, "Models for map building and navigation", *IEEE Trans. on Systems, Man, and Cybernetics*, 23, 1276-1301, 1993.

[17] B. Poucet, "Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms", *Psychological Review*, 100, 163-182, 1993.

[18] A. Redish, "The Hippocampus as a cognitive map II", Computational Models of Neural Systems, 1995.

[19] S. Seed, "Mobile robot map building using landmark recognition in a behavior-based system", Master's Thesis, Machine Intelligence Laboratory, University of Florida, 1994.

[20] K. Snow, "THOMAS - midterm report", Senior Design, Machine Intelligence Laboratory, University of Florida, 1995.

[21] D. Touretzky, A. Redish, H. Wan, "Neural representation of space using sinusoidal arrays", NIPS, 869-884.

[22] C. Vanderwolf, "Hippocampal electrical activity and voluntary movement in the rat", *Electroencephalography and Clinical Neurophysiology*, 26, 407-418, 1969.

[23] J. Winson, "Loss of hippocampal theta rhythm results in spatial memory deficit in the rat", *Science*, 201, 160-163, 1978.

[24] D. Zipser, "Biologically plausible models for place recognition and goal location", *Parallel Distributed Process.*, 2, 432-470, 1986.