IEEE/SOFT International Workshop BIES'95
May 30 - 31, 1995, Tokyo, Japan

Self-Localization in Dynamic Environments

Uwe R. Zimmer

University of Kaiserslautern - Computer Science Department
67653 Kaiserslautern - Germany
Phone: …49 631 205 2624 - Fax: …49 631 205 2803 - Telex: 4 5627 unikl d
e-mail: uzimmer@informatik.uni-kl.de

Self-localization in unknown environments respectively correlation of current and former impressions of the world is an essential ability for most mobile robots. The method, proposed in this article is the construction of a qualitative, topological world model as a basis for self-localization. As a central aspect the reliability regarding error-tolerance and stability will be emphasized. The proposed techniques demand very low constraints for the kind and quality of the employed sensors as well as for the kinematic precision of the utilized mobile platform. Hard real-time constraints can be handled due to the low computational complexity.

The principal discussions are supported by real-world experiments with the mobile robot "ALICE"[1].

Keywords: **artificial neural networks, mobile robots, self-localization, self-organization, world-modelling**

1. Motivation

Based on the idea that a (really useful) mobile robot should be adaptive in its behaviour regarding its current environment, a dynamic world modelling method is introduced in this article. The "really useful" attribute includes sufficient simplicity, reliability, and stability, which are therefore considered basic demands for the proposed approach.

High precision, metric approaches ([2], [5], [8], or [14]) demand very reliably and accurate sensor devices, as well as large computational power. Their applicability on small, and light weighted platforms is limited. During the last few years, "qualitative methods" has been proposed to overcome mainly problems regarding complexity and stability. Works utilizing qualitative modelling for self-localization and navigation include the basic article from Kuipers introducing the term "qualitative map" in [6], the ultrasonic clustering techniques published by Kurz [7], the recently published work of Tani [11] based on local sensor-sequences rather than on explicit topology, and the adaptive, topological models introduced by Prescott [10].

In order to focus on the main issues on this field, the robot's world is designed to be simple, but still of practical relevance. The project as well as the mobile, experimental platform itself will be called *ALICE* in the following. This is not an abbreviation, but just a name.

1-1. The Assumed World

The world of *ALICE* is constructed on the basis of some elementary assumptions. Practical relevance and sufficient simplicity are two contradictory goals that have to be unified in order to define a universe for the present robot. The practical aspects of *ALICE's* world should allow real world tests, which include all relevant problems of mobile robots not having a possibility of (external) global positioning. On the other hand a certain degree of simplicity is needed to allow plausible simulations and to achieve the possibility of generalizing the results to classes of robots and sensors. Furthermore biologically implausible features should be avoided. The central, resulting approach covering all these streams have to be considered regarding certain categories of the world:

- **Position**
  ALICE is allowed to estimate its relative position by dead-reckoning. This is biologically plausible, when the accuracy does not exceed certain limits. Moreover a system being able to handle significant errors in the position measurement contains a wide practical relevance. Another approach would be to dispense with an explicit position at all. Thus the self-localization could be performed on the base of typical sequences of impressions (sensor-readings) as shown in [11].
- **Analog sensor readings**
  ALICE has access to continuously varying features of the outer world, regarding time and movements of the robot itself. According to the biological fact that even the mammal eye is not able to measure absolute physical quantities (in

1. *The project ALICE is supported by the EU-project DG XII, F-5 (Teleman)*

this case light intensities), the sensor models applied for *ALICE* are not calibrated. Moreover the sensor models are non-linear, contain a large degree of noise, and can fail completely.

- **Binary sensor readings**
  Some impressions from the outer world are of a mainly binary nature as for example simple forms of whiskers or other touch sensors. This is also included in *ALICE's* world, in order to demonstrate that spontaneously varying signals can be handled in the proposed world model.

- **Error Tolerance**
  Two kinds of error tolerance have to be discussed. First the error-tolerance in learning or adaptation phases. Here any kind of "miss" is allowed in *ALICE's* world. Due to the fact that the sophisticated mechanics of current robots are not error-tolerance at all in comparison to biological systems, the robot's world itself has to be limited to allow learning in trial-and-error phases. Error-tolerance during the application phase of the system is the second aspect that has to be mentioned. *ALICE* does not distinguish an explicit learning and application phase, i.e. adaptation to the current situation or environment is a lifelong ongoing process in a changing world. This is the first reason for a error-tolerant world at any time. But furthermore the world model as it will be described below includes system-immanent inconsistencies that prevent any globally correct "blind" planning. Any plan may fail in *ALICE's* world and must then be replaced by a more recent and probably better plan. In the opinion of the author, this is not a weakness but an attribute of the real world, which should be included in any robot's universe of practical relevance.

- **Structured Environment**
  ALICE's world should be sufficiently "rich" to enable self-localization at positions or areas of a certain density. That means that the sensor readings must change significantly in travelling distances that are small regarding the relative position corruption. On the other hand the frequency of changing features while moving must be low enough, in order not to generate stochastic signals regarding the robot's sampling frequency. Actually this frequency have to be much lower than given by Shannon's sampling theorem, due to the degree of statistically stability that should be achieved applying fuzzy and unreliable sensor models.

These assumptions about *ALICE's* world lead immediately to the components and structure of *ALICE* (the mobile platform) itself.

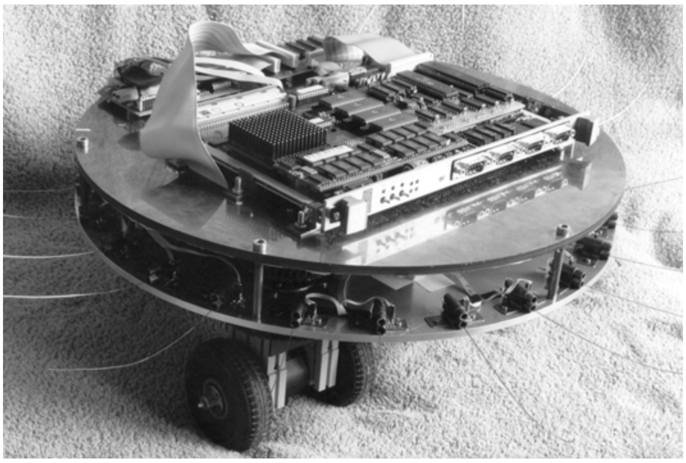

figure 1 : ALICE

1-2. Platform

The physical realisation of *ALICE* and its components will be discussed shortly in this section. For an optical impression of *ALICE* please refer to figure 1. The included features are just sufficient for an autonomous mobile robot operating in realtime in a universe as described in the section above. The platform's kinematics allow an omnidirectional motion control, i.e. *ALICE* is able to go directly towards any direction without any preparatory manoeuvring. The fact that the platform is round and symmetrical regarding the employed sensor devices, skips any movements that would be needed, if certain orientations or perspectives would result in other perceptional or kinematical possibilities. All (six) wheels are driven and steerable. Thus even the propulsion system is completely symmetrical. In order to calculate a relative spatial position the wheel revolutions (measured at the gears) are counted. The achieved accuracy by this dead-reckoning is limited by an (drift-) error of 20 to 25% (depending on the kind of performed movements and the condition of the floor). This position error is "sufficient" to ensure that any weak position correlation technique will fail very soon.

The binary sensor system is realized by 24 simple whiskers mounted symmetrically around the border of *ALICE*. Each whisker is 17 cm long and mounted in a small metal tube, giving a binary signal, whenever the whisker is sufficiently pushed to establish a contact between the whisker and the tube. This "sufficiently" is very fuzzy and depends on the angle, force and speed of the pushing object. Therefore the angular resolution detecting for example a straight wall is below 15° or in some configurations even below 30°.

The sensor system giving continuously varying signals is represented by 24 passive light sensors, measuring the light intensity detected from a certain direction, with an angular resolution of approximately

20°, i.e. every light impressions in a cone of 20° is integrated into one scalar sensor reading. The light sensors (passive photo resistors) are mounted together with the whiskers and are therefore also distributed symmetrically at *ALICE's* border. Please refer to figure 2, giving an impression of the physical realisation of the light sensor - whisker combination. The upper tube hosts a photo resistor, where the lower tube caries the base and the detector for a 17 cm long (steel-) whisker. The tubes are 2.5 cm long.

The on-board computer power is limited to one standard CISC-CPU (currently: MC68040 at 33 MHz equipped with 16 MB random access memory and standard I/O). It is intended (and shown in the following sections) that this computer power is sufficient for a full realtime implementation of all world modelling algorithms. The term "realtime" can be roughly and pragmatically approximated here by the demand that the machine should be able to go continuously at full speed (25 cm/s) where the world model is adapted continuously without any time-jitter or delay.

The author would like to emphasize that the low reliability and resolution of the employed sensor devices together with the low relative position accuracy and the limited computer power is *not* a weakness of the system but is chosen *intentionally*. A low-precision system like *ALICE* is an adequate experimental platform for any world modelling and control technique, which is intended to be stable and reliable in a real world. Any sensor and kinematic system with features superior compared to *ALICE* (i.e. almost any sensor system) promises a further improvement in terms of speed and precision but not with regard of the discussed, principal abilities.

2. Topological World-Modelling

The central motivation of **q**ualitative **t**opological world models (**QT**-Models) is the basic mobile robot task: "Recognize places you have seen before!". In this article this task will be approximated by extracting "situations" (i.e. recognized places) together with their topological neighbourhood from the current sequence of sensor-samples, rather than modelling the boundaries of the detected obstacles and objects in a metric manner. Assuming a stable situation-recognition-process and a technique for moving between distinct situations, the concept of a qualitative, topological world model suggests a human-motivated basis for a navigation. The main concept has already been proposed by Kuipers et al. [6], but here the construction process was carried out using explicit rules, not statistical techniques. Therefore the real-world abilities of the Kuipers approach are, in the opinion of the author, limited.

The world model proposed in this article is based on clustering techniques introduced by Kohonen ("self-organizing-maps", [4]) and Fritzke ("growing cell structures", [1]) together with some previously proposed extensions by this research group [3]. Due to a couple of specific autonomous robots-constraints, these structures are modified to cope with realtime-aspects, lifelong learning, "local forgetting", and correlation.

To make the term "situation" in this context more precise, figure 3 shows a typical situation generated in *ALICE's* test-runs. The inner circle describes the distribution of light-impressions received from particular directions, the outer circle shows the smoothed touch-values originate from contacts with obstacles in various directions. The two-number pair in the centre gives a rough approximation of the geometric position of this "situation".

3. Methods

This section will introduce the technical details of the proposed topological world model. The algorithms following are expressed in general terms only ignoring computational details.

3-1. Pre-Processing

Following the idea of representing situations, consisting of readings from different kinds of sensors in a way that they can be compared in one step, and by

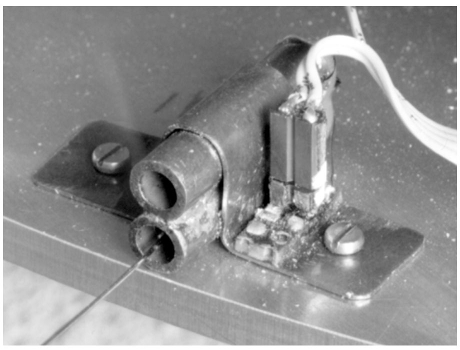

figure 2 : Light sensor - whisker combination

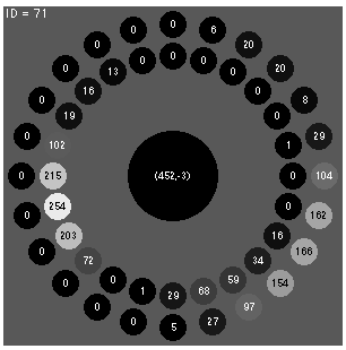

figure 3 : A "situation"

employing a simple norm, the sensor-samples have to be preprocessed to form a vector of unified elements. In the current system passive light and tactile sensors as well as an *(x, y)*-position produced by odometry are available. Considering the fact, that the angular resolution of the tactile sensors is very low, each vector of tactile readings is smoothed by applying gaussian functions. Finally the sensor samples from different types of sensors are weighted and concatenated to produce a "situation-vector", or more briefly a "situation" (consisting of 50 values in the example given).

In the following, sensor situations will be indicated as S; the position or x, y-part of such situations as p.

3-2. Adaptation

As a basis for the network model, the euclidian norm is applied to calculate distances between sensor situations, d, and distances between positions, g, respectively.

Consider a network N consisting of a number of cells c_i, which are connected with respect to the topological neighbourhood of the situations $S(c_i)$ attached to each cell. Then at each adaptation step the cell c_{opt} with the smallest situation-distance d_{opt} to the new input situation S_x is determined according to:

$$\forall c_i \in N: d_{opt} = d(S(c_{opt}), S_x) \le d(S(c_i), S_x) \quad (1)$$

In order to limit the effort for this adaptation to a constant the search area is limited by the geometric distance g_{search}. In the current system this is done by applying adequate data-structures to the network-management. The selected cell c_{opt} and all its topological neighbours are then adapted according to:

$$c_{opt}^{new} = c_{opt} - (\varepsilon_o \cdot d_{opt}) \quad (2)$$

$$\forall c_n^j \mid a(c_n^j, c_{opt}) > 0: \; c_{n_j}^{new} = c_{n_j} - (\varepsilon_n \cdot d_{n_j}) \quad (3)$$

where $a(\cdot, \cdot)$ is the adjacency-function of the network. The "classification error" d_{opt} is then added to a total classification error d_{total} attached to the cell c_{opt}.

$\forall c_i \in N$ (after n adaptation steps):

$$d_{total_i} = \sum_{t=1}^{n} hit_{i,t} \cdot d_{opt_t} \quad (4)$$

$$\text{where } hit_{i,t} = \begin{cases} 1 \; ; c_i = c_{opt_t} \\ 0 \; ; c_i \ne c_{opt_t} \end{cases} \quad (5)$$

In order to decrease the adaptation speed of a well adapted network, the parameters ε_o and ε_n are controlled by:

$$(\varepsilon_o^{new}, \varepsilon_n^{new}) = \begin{cases} (\varepsilon_o \cdot \varepsilon_\Delta, 0) \; ; (d_{opt} \le d_{acc}) \\ (\varepsilon_o^{init}, \varepsilon_n^{init}) \; ; (d_{opt} > d_{acc}) \end{cases} \quad (6)$$

where: $0 < \varepsilon_\Delta < 1$

and ε_o^{init} resp. ε_n^{init} are the initial values of the parameters ε_o and ε_n.

In each adaptation step, where d_{opt} is larger than a tolerated error d_{acc}, a global counter n_{miss} is incremented. This counter will be used as a measurement for the need for change in the network structure. An update-counter u_{opt} attached to c_{opt} is incremented and will be used as an indicator for the stability of this specific cell.

In order to use the high speed of this adaptation process to achieve better adaptation, each situation is presented several times k to the network. A constant delay of l sensor-sample-time-slots before the current sensor situation affects the network is also found useful (see section "Correlation" below). Accordingly a learning set holding $(k \cdot (l+1)) - 1$ situations is implemented.

3-3. Growing & Shrinking

At start-up time of the system, there are no cells; the network is empty. So the common problem finding a "good" initial state of the network is avoided, but there is a need for some growing strategy. In the present system, two growing strategies are applied. The first is called "**spontaneous insertion**", the second "**statistical insertion**". In the first, new cells, representing the current sensor situation, are inserted when the distance between the current sensor situation and c_{opt} exceeds a certain limit ρ_s (in the special case of an empty network this strategy produces the first cell). In the second strategy a new cell is inserted in the middle between the cell with the highest "degree of movement" c_{runner} (measured by the cell attribute d_{total}) and its farthest topological neighbour c_{far} every n_{insert} "miss-classifications" (measured by the global counter n_{miss}). The new cell is instantiated with mean-values of c_{runner} and c_{far} for position and light-intensity, but with minimal values for touch-information.

Another aspect of growing relates to the topological connections between cells. Assuming that c_{opt} has just changed from c_{opt}^{old} to c_{opt}^{new} in two consecutive adaptation steps, and that the cell c_{opt}^{new} has m other neighbours c_j ($a(c_{opt}^{new}, c_j) > 0$), the following changes in connection weights are imposed:

$$a(c_{opt}^{new}, c_{opt}^{old}) = 1 \quad (7)$$

$\forall j, (1 \le j \le m)$:

$$a^{new}(c_{opt}^{new}, c_j) = a(c_{opt}^{new}, c_j) - (a_{red}/m) \quad (8)$$

with $0 < a_{red} \le 1$

A connection with a weight ≤ 0 is regarded as non existent. Thus the deletion of cells is now straight forward. A cell or a cell-cluster with no connection is removed.

3-4. Correlation

Three degrees of freedom out of the internal representation ((x, y)-position and orientation) will be corrupted by drift effects or other errors, if they are not continuously correlated to the world model. Even a stable and error-tolerant network structure will not be able to produce stable world models, if the robot's position is corrupted. The obvious approach is to correlate the internal position continuously with the world model built up so far. On the other hand the robot's position is integrated in the world model. This mutual stabilizing technique is very useful applied to local world models. But in fact each representations are updated with the faults, errors and noise from the other. This principal problem prevents a globally consistent world model of arbitrary size including position when only local information is available.

One strategy, used to stabilize the internal position, is to delay the integration of the current sensor situation S_c by a number of adaptation steps l. In this way the internal position is correlated on the basis of the formerly integrated world-knowledge at this point, and not on the base of S_c, which would not make any sense. An estimation p_{est} of the current position is produced using the sensor-distances d_i (to S_c) and the number of updates u_i of all cells c_i in the immediate topological neighbourhood of S_c. The estimates produced in this way, are not necessarily near to the "correct" position in every case. Thus the internal position is not changed to p_{est}, but is only "moved" towards this point by a fraction ε_{pos}. The fact that these estimation errors are normally distributed in the long term, together with the fact that the remaining position-errors are integrated into the world model after a certain delay, leads to a stable behaviour of the local world model.

The third degree of freedom, the orientation of the robot, is also corrupted by drift and other effects. In order to correlate orientation with the internal representation and the world model, it must be assumed that the current position is correct to within a certain tolerance. Based on the geometric distance, the nearest cell c_{near} is determined by:

$$\forall c_i \in N: g(S(c_{near}), S_x) \leq g(S(c_i), S_x) \qquad (9)$$

To obtain a comparable measurement for the orientation based on sensor impressions including position, we interpret the sensor values as polar coordinates:

$$l_i \Rightarrow (l_i \cdot sin(i \cdot \frac{360°}{s}), l_i \cdot cos(i \cdot \frac{360°}{s})) \qquad (10)$$

$$t_i \Rightarrow (t_i \cdot sin(i \cdot \frac{360°}{s}), t_i \cdot cos(i \cdot \frac{360°}{s})) \qquad (11)$$

where l_i are the light-readings, t_i are the tactile readings and $0 \leq i \leq s - 1$. A possible derived measurement for the orientation (modulo 180°) is a linear regression O on the $2s$ polar "sensor-points". The angular difference between the orientation of the regression of the current sensor scan O_c, and the orientation of the regression O_{near} of the nearest cell c_{near}, is interpreted as the current orientation-error. As in the case of the position correlation above, the internal orientation is only "moved" by a fraction ε_{or} towards the estimated orientation. In the current system only the light sensors are employed to determine the orientation estimates. The internal orientation is continuous and not quantised according to the angular sensor-resolution. Thus the incoming sensor-readings are interpolated linearly and rotated according to the current orientation of the robot.

Due to the fact that the remaining position-drift is integrated with the world-model, a global drift of the whole model may occur. This is not necessarily a problem if the world model is drifting as a whole, but as the world model expands and some areas are visited in a sporadic manner, different areas will drift in different ways, resulting in an inconsistent world model. Therefore a cell-specific slow-down of the adaptation speed is introduced when cells are "confirmed" at least n_{fix} times (i.e. $u_{opt} \geq n_{fix}$).

4. Experiment

In this section the author tries to emphasize the real world aspect of the *ALICE* project, i.e. the world modelling should be stable regarding the assumed world introduced above. The behaviour of *ALICE* is documented under a couple of critical conditions, like inadequate parameters, certain sensor weights, lost correlation, and dynamic environments. The most important, observed feature during the test-runs is the internal position of *ALICE*. Whenever *ALICE* loses its position respectively orientation correlation, the world model is obviously useless for the intended purpose. Furthermore some heuristics regarding the quality of the network can be applied and evaluated (in case of a stable internal position):

- **Generalization**
  The number of generated cells in the network should be small. Thus requirements regarding realtime, memory capacities, and generalization can be fulfilled. Moreover only a sufficiently low density of cells enables the correlation abilities.
- **Topology**
  The generated network should have a topological equivalent in the real world. Moreover only a low degree of interceptions between the connections should be accepted. Otherwise the graph search methods of the navigator's planning component will fail.
- **Consistency**
  The cells as well as the generated connections

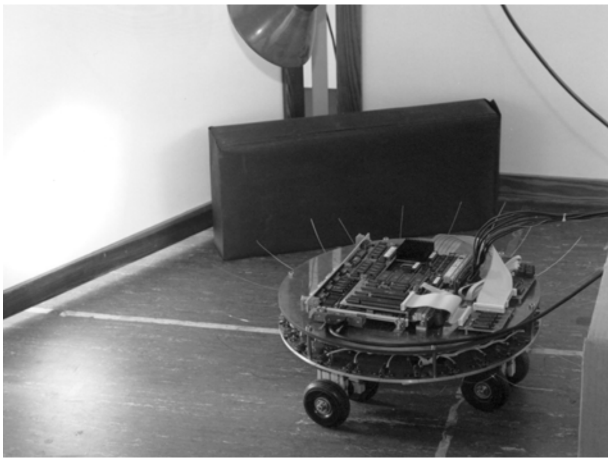

figure 4 : Test scene

should remain in areas, reachable for *ALICE*. The adaptation of the network in the form of representation through mean values given no general guarantee for this feature. Situations represented by the cells in the network are not necessarily observable or existing in the real world. Especially the neighbourhood adaptation plays a critical role in the generation of inconsistent network areas.

The tests are performed in a polygonal (roughly rectangular) environment of approximately three times three metres with a round pillar (diameter: 45 centimetres) slightly asymmetrical in the centre. The borders are straight wall segments, each of a length of one metre. Two respectively three light sources are mounted at the border of the test environment producing any kind of reflections and shadows. Depending on the time of the day the sunlight gives a fourth light source, with completely other characteristics than the spots with their artificial light. Especially the continuous change of the sunlight makes the test environment dynamic without further engagement. Figure 4 gives an impression of the actual real world environment. The bright spots at the left are light reflections.

The shown results depend critically from the strategy and order of gathering sensor readings from the actual environment. The set of applied strategies (called "exploration") for these experiments, is discussed in [13].

4-1. Critical Network-Parameters

As one representative out of the group of network parameters, the critical influence of the parameter ε_o is shown. ε_o controls the adaptation speed of the cells (ε_n is changed proportional to ε_o in this test, i.e. $\varepsilon_o = 10 \cdot \varepsilon_n$). First, ε_o is set to a large value (three times larger then the usual value of 0.03). The resulting world modelling performed by *ALICE* shows a

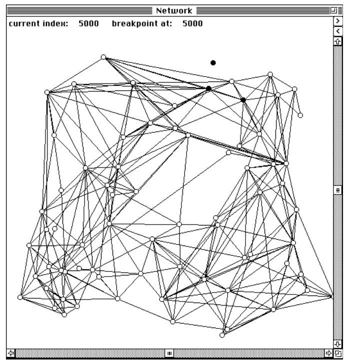

Figure 5a: $\varepsilon_o = 0.1$

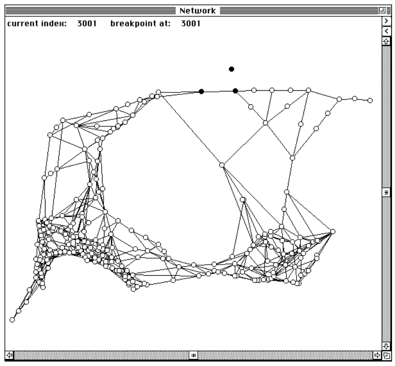

Figure 5b: $\varepsilon_o = 0.005$

very "lively" behaviour (figure 5a). The cells are moving quickly and only a small number of cells is required to represent the whole environment. But the connections are too muddled to be useful for a navigator and the representation is not consistent for example regarding the pillar in the centre. The counter test of reducing the adaptation speed of the network to one sixth of the usual value results in a even worse situation. A huge amount of cells is inserted, due to the fact that the slowly moving cells cannot integrate the quickly changing, sampled situations. Even the position and orientation correlation fails, when the test as shown in figure 5b is continued.

4-2. Sensor Weights

The influence of the sensor weight concerns mainly the kind of produced map. The influence of light and touch information to the sensor situation is given by the distribution and strength of these readings in a typical working environment, but the weight for the position has moreover an additional meaning. This weight controls (beside other parameters) the topologic equivalence of the network and real world regarding the robot's spatial position. If the position weight is chosen to small, the topologic neighbourhoods in the network are controlled by the light and touch information only and the test results in wide, intersected connections, without any geometric equivalence. On the other hand, if the meaning of the

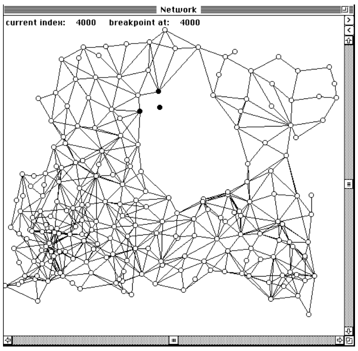

figure 6 : large position weight

position for the sensor situation is set too high (figure 6), the geometric equivalence seems perfect, but the number of cells needs to be rather high, and even worse, the position correlation is getting very critical. Due to the fact that the internal position is mostly correlated on the base of other (network-) positions (i.e. not on the base of complete sensor impressions) this world model must fail after a certain time.

The actual setting of the position weight depends on the robot's dead-reckoning precision as well as on the degree of needed geometric consistency.

4-3. Correlation

The importance of the position- and orientation correlation for the world modelling (and of course for the navigator and explorator) can be shown by two test-runs, where one of the correlation processes is deactivated in each case. The results are shown in figure 7 for a test-run without orientation correlation. Due to the fact, that these experiments without correlation are completely instable, the shown qualitative topologic maps are arbitrary but representative. None of these test-runs could be completed without a total loss of consistency in the world model. This is not surprising, with *ALICE's* position drift-error of 20 to 25% and rotational drift of up to 0.5°/metre in mind. But the principle loss of consistency does *not* depend from the actual drift rate – as long as there is any drift at all (as in all real world systems) the effect will occur sooner or later.

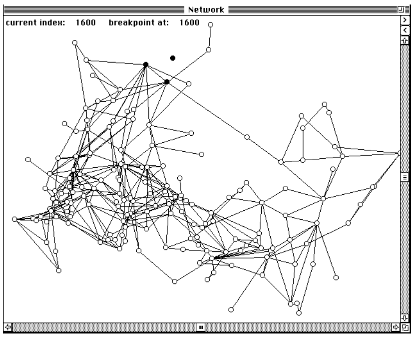

figure 7 : …without orientation correlation

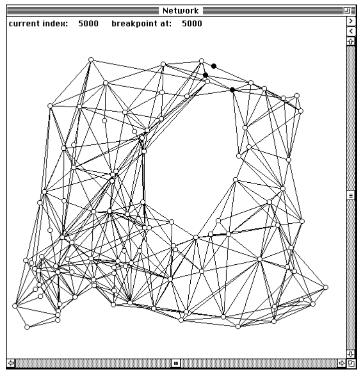

figure 8 : static, 5000 samples

4-4. Static and Dynamic Scenes

The final section of the experiments will discuss the development of working qualitative, topological world models in static and dynamic environments. The world model in a static environment reaches an equilibrium state after sufficient exploration of all available features. *ALICE* needs approximately 15 to 20 minutes to build up a QT-map of the static test environment as demonstrated in figure 8. The final state (after gathering 5000 sensor situations) represents the geometric features, the light distribution as well as a network graph well suited for the navigator. As introduced in section 3, the network holds much more information than shown in these figures. The cells contain whole sensor situations together with statistical values about their history, where the connections are attributed by degrees of confidence. All

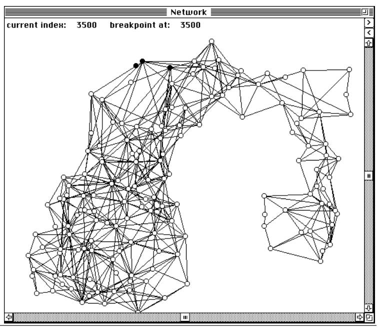

Figure 9a: dynamic, 3500 samples

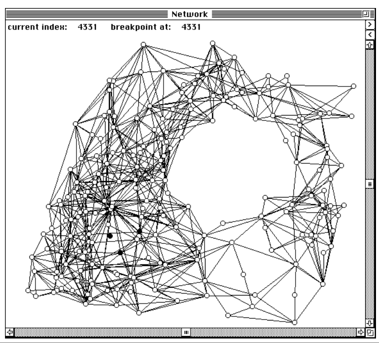

Figure 9b: dynamic, 4331 samples

this information can be employed by the planning or execution (driving) phase of any navigator.

The second world model developing example shown in figure 9a and figure 9b demonstrates the abilities to adapt to a changing world. In order to make the effect obvious, the formerly closed circle in the environment is cut off at the lower end until the gathering of sensor situation 3500 (figure 9a). Up to this situation, the world model shows a clear gap in the lower part. Although the absolute position error between the two sides of the gap is larger than it would be without the splitting wall, then gap is closed smoothly after another 1500 training steps (and of course after removing the introduced obstacle). Due to the careful and smooth removal routines, the two worlds coexists for a certain time, until the world with the gap is completely “washed-out”. The navigator may take advantage from the fact that the confidence values of the connections distinguish between most recent and established information.

5. Conclusion

As the reader may have expected, the choice of a world model will depend significantly on the task, but some guidelines may be derived from the discussion. One aspect is reliability, for example in a non-error-tolerant environment. If that is to be a central aspect of the robot-task, an exact model may be required to be able to plan safe paths. Similarly, if a guarantee of accuracy when following a path is needed, the exact geometric information may be necessary.

On the other hand, if the main focus is on simplicity, stability or qualitative aspects of the task, the qualitative topological map techniques may be the first choice. Especially the small requirements for computational effort and sensor equipment together with a high degree of robustness is an unique feature. The experiments have shown real world abilities offering sufficient information for navigation purposes as well as a stable self-localization method.

References

[1] Bernd Fritzke
Growing Cell Structures - A Self-organizing Network for Unsupervised and Supervised Learning
Technical Report 93-026, International Computer Science Institute, Berkeley, California

[2] Peter Hoppen, Thomas Knieriemen, Ewald von Puttkamer
Laser-Radar based Mapping and Navigation for an Autonomous Mobile Robot
IEEE International Conference on Robotics and Automation 1990, Cincinnati, Ohio, May 13-18, 1990, pp. 948-953

[3] Herman Keuchel, Ewald von Puttkamer, Uwe R. Zimmer
SPIN - Learning and Forgetting Surface Classifications with Dynamic Neural Networks
Proceedings of the ICANN '93, Amsterdam, The Netherlands

[4] Teuvo Kohonen
Statistical Pattern Recognition Revisited
Advanced Neural Computers / R. Eckmiller (Editor), Elsevier Science Publishers B.V. (North-Holland), 1990

[5] D. Kuan, J.C. Zamiska, R.A. Brooks
Natural Decomposition of Free Space for Path Planning
IEEE International Conference on Robotics and Automation, St. Louis, Missouri, 1985, pp. 168-173

[6] Benjamin J. Kuipers, Yung-Tai Byun
A Robust, Qualitative Method for Robot Spatial Learning
Proceedings of the AAAI 1988

[7] A. Kurz
Building Maps based on the Learned Classification of Ultrasonic Range Data
Intelligent Autonomous Vehicles (IAV) '93, International workshop, University of Southampton, Hampshire, United Kingdom, 18-21 April '93, pp. 193-198

[8] T. Lozano-Pérez
Spatial Planning: A configuration Space Approach
IEEE Transactions on Computers, Vol. C-32, No. 2, 1983, pp. 108-120

[9] Maja J. Mataric
Integration of Representation Into Goal-Driven Behaviour-Based Robots
IEEE Transactions on Robotics and Automation, Vol. 8, No. 3, June '92

[10] Anthony Prescott
Explorations in Reinforcement and Model-based Learning
PhD-thesis, Department of Psychology, University of Sheffield, GB, 1994

[11] Jun Tani, Naohiro Fukumura
Learning Goal-Directed Sensory-Based Navigation of a Mobile Robot
Neural Networks, Vol. 7, No. 3, pp. 553-563, 1994

[12] Gerhard Weiß, Christopher Wetzler, Ewald von Puttkamer
Keeping Track of Position and Orientation of Moving Indoor Systems by Correlation of Range-Finder Scans
IEEE International Conference on Intelligent Robots and Systems '94 (IROS '94), September 12-16, 1994, Munich, Germany

[13] Uwe R. Zimmer
Robust World-Modelling and Navigation in a Real World
submitted for publication (Neurocomputing, special issue), 1995

[14] Uwe R. Zimmer
Connectionist Decision Systems for a Visual Search Problem
Proc. of the IIZUKA ´94, Fukuoka, Japan, August 1-7, 1994, Invited paper

[15] Uwe R. Zimmer & Ewald von Puttkamer
Realtime-learning on an Autonomous Mobile Robot with Neural Networks
Proc. of the IEEE Euromicro '94 Realtime-Workshop, Västerås, Sweden, June 15-17, '94