H. Hügli, G. Maître, F. Tièche, C. Facchinetti "Vision-based behaviours for robot navigation", Proceedings of the Fourth Annual SGAICO Meeting, Neuchâtel, Septembre 1992

VISION-BASED BEHAVIOURS FOR AUTONOMOUS MOBILE ROBOTS

Heinz Hügli, Gilbert Maître, Claudio Facchinetti, François Tièche
University of Neuchâtel, Institute for Microtechnology
Rue Tivoli 28, CH-2003 Neuchâtel, Switzerland

ABSTRACT

This paper is a contribution to the investigation of the behavioural approach in robot vision. It describes the principle and implementation of a behaviour-based vision architecture for autonomous mobile robots. The paper gives some motivations for the behavioural approach and presents the selected vision architecture. The various behaviours selected to allows basic capabilities for vision-based navigation are: *going towards*, *going to, going along, obstacle avoidance, obstacle detection, landmark following and free-space mapping*. They are described together with the two vision systems they rely on: vision by structured light and active vision by landmarks.

1. Introduction

The traditional AI approach to design intelligent systems states that intelligence operates on a system of symbols. In perception based systems, this approach has not brought the expected success. Among the reasons that can be evoked for this momentary failure, there is the big difficulty to establish a relevant link connecting perception and symbols. Because perception is an active and task dependent process, we observe that such links change in time and from context to context, and that such changes make many symbol interpretation meaningless.

In contrast, the behavioural approach states that to build a system that is intelligent, it is necessary to have its representations grounded in the physical world. It tends therefore to use the physical environment of the intelligent system as the only representation. This means building a system bottom up and making high level abstractions concrete. Also, goals and desires have to be expressed in terms of physical, measurable actions [Broo90].

The traditional AI and the behavioural approaches lead to different methodologies for building mobile robots [Broo86]. The traditional AI approach decomposes the intelligence in functional processing blocks whose combination provides overall system behaviour. To get any behaviour at all, it is necessary to combine together many modules; improvement of the system proceeds by improving the individual function modules.

In contrast, the behavioural approach bases its intelligence on individual behaviours whose coexistence and co-operation let more complex behaviours emerge. Here, each module generates behaviour and improvement of the competence of the system is obtained by adding new modules to the system.

In the following , we will present a vision system for an autonomous mobile robot that was developed according to the behavioural approach. In section 2, we describe behaviours and show their use in the robot vision system. The vision system itself and the various vision-based behaviours are then described in detail in the two next sections. Finally, the current implementation and projected further work is described in section 5.

2. Behaviours that solve the robot's task

According to the behavioural approach, a robot task results from a sequence of stereotyped actions we call behaviours. The set of behaviours a robot can perform characterises itself in terms of its capabilities to interact with the environment. The behaviour itself is an action that is started by a stimulus and that exists as long as this stimulus exists.

Animal behaviours

The behavioural approach is inspired to some extend by the animal world in which such elementary behaviours can be observed [Mori88]. As an example, the bee, despite a poor visual system, easily navigates to and from its hive. To reach its hive, the bee could for instance proceed in two steps: first, follow a characteristic fence, and, as soon as the hive is visible, go toward it. The bee's goal is reached in two successive behavioural steps. For now, we consider the behaviour to be related to a specific input and to offer a partial navigation solution.

In the *going along* behaviour of the bee, it is the visible linear structure of the fence that is the sign pattern that triggers and maintains the behaviour. In the *going to* behaviour that leads the bee to its hive, the red spot of the hive acts as the sign pattern.

Behavioural architecture

In general, as several stimuli can be active at the same time, several behaviours may become active simultaneously. Fully independent behaviours will then run concurrently while behaviours which share some common resources are incompatible and exclude each other. Among several behaviours competing for a common resource, a single one can be selected. The selection is performed according to a decision scheme dictated by the planner.

In the behavioural context, planning acts on the system by allowing behaviours to run or not, and by acting on the decision scheme used for the selection of one among the competing behaviours. It does so in a context sensitive manner.

A direct advantage of the behavioural approach is that it permits to develop the behaviours as quite independent units. Apart from some interaction constraints, the developer does not have to care about the implementation of other behaviours. Each of these is a separate process which activation merely depends on the presence or absence of a stimulus.

We make a distinction between an external and an internal behaviour. An external behaviour is one that performs an action within a feedback loop across the robot environment whereas the internal behaviour does not act in an explicit loop; we call internal behaviour one that builds or updates internal knowledge from perceived stimuli.

The independence of behaviours as well as their modularity make them suited for reuse as complete behavioural units. By combining several such units, other, more sophisticated behaviours can be built. It is believed that behavioural abstraction can be obtained this way.

Robot task and environment

The mobile robot is located in a building, moving horizontally on flat ground. Walls and different obstacles make up its environment. The robot must fulfil a task. Among the planned tasks there is the task to stow chairs in a room. The robot is further autonomous, which means that it is capable to perform its task in an unpredictable environment.

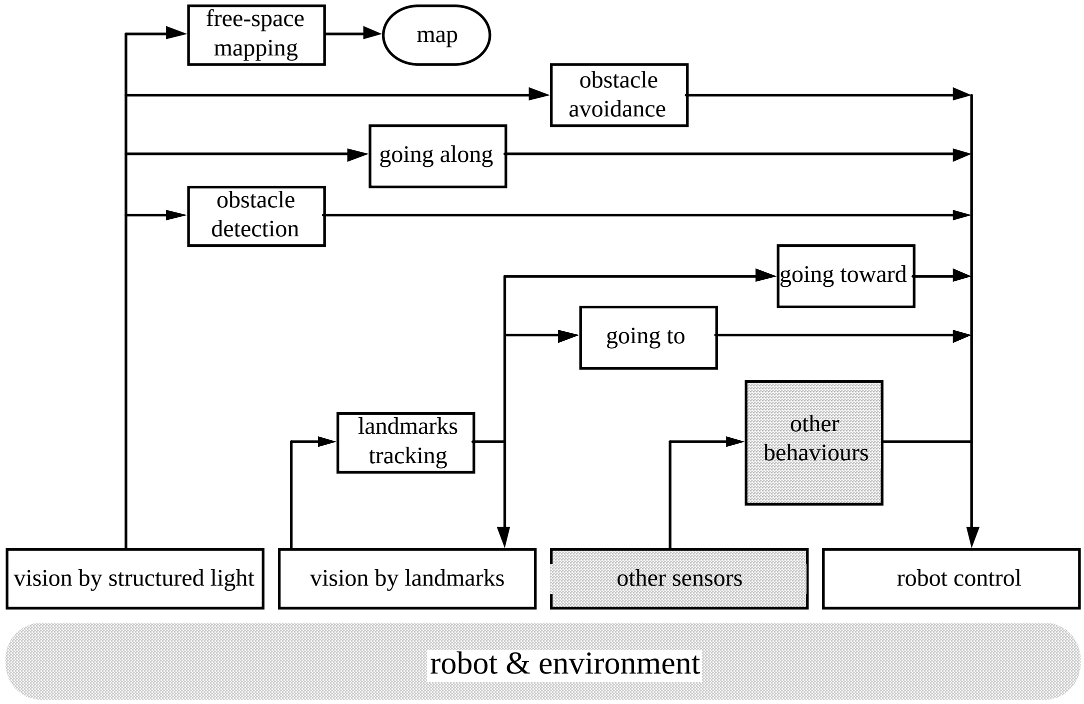

Fig. 1 Architecture of the behavioural vision system

Vision-based behaviours

Vision-based behaviours are characterised by the fact that their stimulus is a visual pattern called sign pattern. It is this sign pattern that triggers the behaviour and that maintains the behaviour active as long as it exists.

The design of our vision system for autonomous mobile robots was inspired by several sources as [Mori88], [Asad90], [THKS88], [Turk88], [Khat86] a. o. Its current architecture is shown in figure 1 where we recognise the following vision-based behaviours: *going towards, going to, going along, obstacle avoidance, obstacle detection, landmark following and mapping*. This set of behaviours allows basic capabilities for vision-based navigation.

The behaviours found in this basic set are of several types. Most of them act in a feedback loop across the environment and fall therefore in the category of external behaviours according to our definition. There is only one internal behaviour –*free-space mapping*– which builds up an internal representation of the environment. Notice also that not all behaviours have the same level of interaction with the environment. For instance *landmark tracking* which acts directly in a loop across the sensing system, directly contributes to the two behaviours *going to* and *going towards*: its result is reused in these two behaviours. The reuse of behaviours builds a stucture in from of a hierarchy.

3. Vision by structured light

A first set of behaviours is based on the measurement of the geometry in front of the robot by a structured light measurement device: *obstacle detection, obstacle avoidance, going along and mapping*.

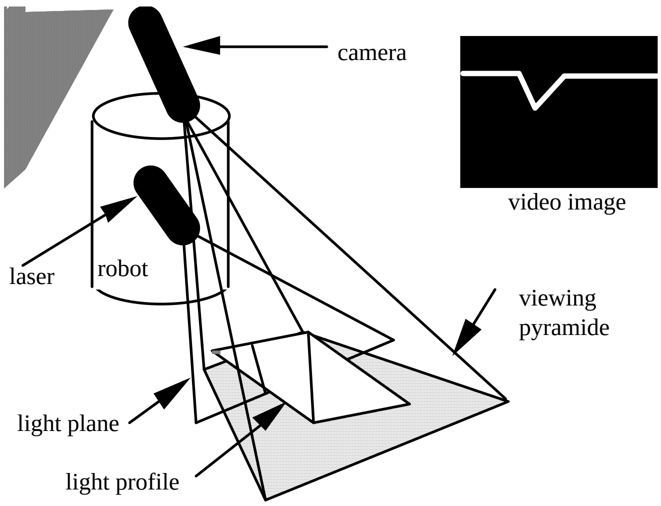

Fig. 2 Vision by structured light

Structured light measurement device

The robot is equipped with a measurement device using the triangulation principle by structured light in order to perceive the environment in front of the robot (fig. 2). The three main components thereof are the laser beam, the video camera and the image processor [HuMa89]. The laser projects a plan of laser light towards the ground, one meter in front of the robot, which intersection produces a stripe of light observed by the video camera.

The stripe of light observed by the camera is a basic straight line in presence of a flat ground and shows a modified profile in presence of obstacles. The image processor analyses the light profile and produces the sequence of line segments which fall on the basic straight line. Formally, the system produces the list s1, s2,… of these segments –ground segments from now– expressed in the robot reference system.

This first device belongs to the class of range measuring devices. There are two important reasons which make the use of such devices very attractive for robot vision. Firstly, they measure geometry, i.e. an intrinsic property of the environment. Secondly, recent developments led to efficient range measuring devices [Ber92].

Obstacle detection behaviour

Obstacle detection is an external behaviour that triggers the obstacle alarm in the presence of obstacles in front of the robot. Its stimulus results from a given configuration of the ground segments observed in front of the robot. Various possible rules apply for the choice of configurations which are valid stimuli for the behaviour. In essence, all use the principle of a verification of the presence of ground segments in the direction of the next predicted robot displacement.

Obstacle avoidance behaviour

Obstacle avoidance is built around obstacle detection and is triggered by it. It analyses the ground segment configuration for a possible next movement. If at least one is found, a new movement command is issued, otherwise a stop command is given.

Going along behaviour

Going along is an external behaviour that keeps the robot moving in the environment along a given linear extended sign pattern. Candidate sign patterns for our measuring device by structured light are curvilinear 3D shapes like pipes, edges, … . Detecting the wall-ground edge leads to a *wall following* behaviour; detecting the depth discontinuity found between a road and a sidewalk leads to a *sidewalk-following* behaviour.

Therefore, the basic mechanism used in *going along* is the detection of the specified sign pattern in the list of ground elements, followed by its tracking. The remaining of the behaviour action is to control the robot's movement to obtain the desired movement along the extended sign pattern.

Free-space mapping behaviour

Free-space mapping is the internal behaviour that keeps track of the ground segments observed in the past. It build a map that integrates the learned knowledge about ground-level space, i.e. space that is free for navigation. This is done successively by reporting the observed ground segments into the map, at the position indicated by the odometry. Interpolating between successive segments leads to a map which displays areas that can be considered free for navigation.

4. Vision by back-reflecting landmarks

A second set of behaviours is based on a vision system by landmarks: *landmark following, going towards, going to*.

Vision device

A fixed video camera is mounted on top of the robot, together with a light source. The light illuminates the scene (fig. 3). Because of their back-reflecting property, the landmarks reflect most light towards the light source, i.e. the camera, where they appear as bright spots on a darker background. This active vision device produces an image of the landmarks found in front of the robot [Fac92].

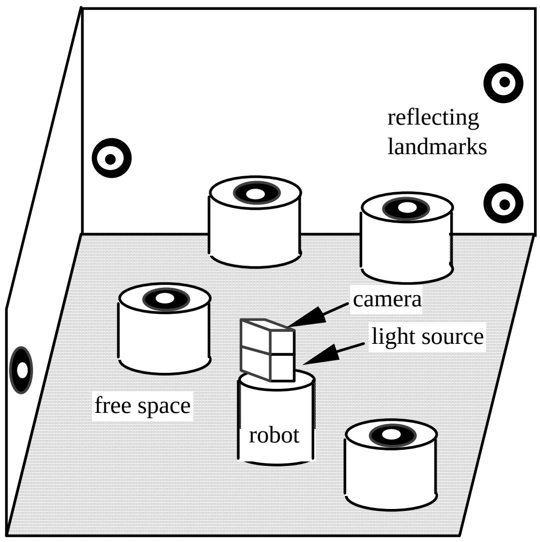

Fig. 3 Active vision of landmarks

***Landmark tracking* behaviour**

This is the behaviour which analyses the scene for detection of landmarks and for keeping track of their position. The landmark acts as the sign pattern for that behaviour and the time sequence of screen co-ordinates is the knowledge extracted from it.

In essence, this is the classical problem of multiple target tracking. The success of the solution is tightly related to real-time constraints. The processing capabilities must be fast enough to keep track with the continuity constraints of the relative movement of the landmarks .

In our current implementation, we have single and multiple landmark tracking at a maximum frequency of 16 Hz. For each landmark, there is a trajectory prediction based on a uniform accelerated movement in the image space.

***Going towards* behaviour**

This external behaviour aims at moving the robot towards a given landmark. It uses the *landmark tracking* behaviour, already described, that tracks a landmark and returns its current position. In *going towards*, this returned position of a given landmark is thus used to move the robot in the direction of the landmark. The difference between current robot bearing and landmark bearing controls the robot movement to the landmark.

***Going to* behaviour**

This is an external behaviour built around *going towards*. It aims at moving the robot towards a landmark but unlike *going towards*, it does not lead the robot straight to the goal; it rather moves it on a longer way. The advantage of this behaviour is that during its longer path to the landmark, the robot sees the landmark under different bearings, a fact which can be used, together with odometry, to determine the distance to the landmark.

In a preferred implementation, the robot starts its behaviour by a paraxial movement with respect to the landmark, a move during which it obtains a first estimate of the distance to the landmark. Then, there is either a move towards the landmark or an additional paraxial movement, depending on the value and confidence of the distance found. This step is repeated recursively until the goal is reached.

5. Implementation and experiments

The vision architecture previously described has been implemented as part of the MARS autonomous mobile robot, which, beside vision-based behaviours, used also other sensor-based behaviours described elsewhere [GaMu91]. MARS

is a Hero 2000 robot controlled by a behavioural architecture implemented on a series of eight Macintosh computers. The computers are interconnected and form a multiprocessing environment.
The current implementation of the vision system offers *going along* in the form of *wall following*, *obstacle detection, single landmark tracking, going towards* and *free-space mapping*.
Several experiments conducted with the robot demonstrated the expected capabilities of the different behaviours both in their separate and combined use.

6. Conclusions

In this paper, we presented a vision architecture that adopts the behavioural approach. The seven behaviours developed so far perform as expected and could be integrated in a common architecture that combines them, some of them in a hierarchy. Simple sequences of different behaviours could be obtained practically. It is now planned to further extend this approach in two direction. The most intriguing is the extension of the system towards more abstract behaviours. The other is the implementation of a second-generation experimental environment.

7. Acknowledgements

This work has been supported by the Swiss national research program "Artificial Intelligence and Robotics" (PNR23) under project number 4023-027037. It is a collaboration with the AI Group, Institute of Informatics, University of Neuchâtel.

8. References

[Asad90] M. Asada, "Map building for a mobile robot from sensory data", IEEE Trans. Systems, Man, Cybernetics, Vol 37, No. 6, Nov./Dec. 1990

[Ber92] Beraldin J.A. & al.," Registered intensity and range imaging at 10 mega-samples per second, Optical Engineering, Jan. 1992, Vol. 31, No. 1

[Broo86] R.A. Brooks, "A robust layered control system for a mobile robot", IEEE Journal of Robotics and Automation, Vol RA-2, no. 1, March 1986

[Broo90] R.A. Brooks, "Elephants don't play chess", Designing autonomous agents, Ed. Pattie Maes, MIT Elsevier, 1990

[Fac92] Facchinetti Claudio & Hügli Heinz, "Two vision-based behaviours for autonomous mobile robots", Proc. Conference on Pattern Recognition and Vision, Lausanne, Switzerland, Jan. 1992

[GaMu91] Y. Gat & J.P. Muller, "Simple World Modelling for Reactive

Navigation", AAAI Fall Symposium, Nov. 1991

[HuMa89] H. Hügli & G. Maître, "3D by structured light: implementation and evaluation of a vision system for small parts", in Optical 3-D Measurement Techniques, Ed. A. Gruen & H. Kahmen, Wien, Wichmann, 1989, pp. 468-477

[Khat86] O. Khatib, "Real-Time obstacle avoidance for manipulators and mobile robots", The international Journal of robotics research, Vol. 5, No. 1, 1986

[Mori88] H. Mori & al., "A mobile robot strategy applied to Harunabu-4", Int. Conf. Pattern Recognition, Rome 1988

[THKS88] C. Thorpe, M. Hebert, T. Kanade, S.A. Shafer, " Vision and navigation for Carnegie-Mellon Navlab", IEEE Trans. PAMI, Vol 10, No. 3, May 1988

[Turk88] Turk M.A. & al., "A vision system for autonomous land vehicle navigation", IEEE Trans. PAMI, Vol 10, Nr. 2, May 1988