

A Control Architecture for an Autonomous Mobile Robot

Maria C. Neves

Superior Institute of Engineering
Polytechnic Institute of Porto
R. S. Tome
4800 Porto
cneves@fe.up.pt

Eugénio Oliveira

Faculty of Engineering
University of Porto
R. dos Bragas
4000 Porto
eco@fe.up.pt

Abstract

In this work we propose a behaviour-based architecture, which can also be seen as a multi-agent architecture, to control a situated and embodied autonomous mobile robot, under a close sensing action relationship. Pre-existing knowledge can play a meta-level role, being used in order to provide an adviser capability which is combined with the basic reactive feature of the robot. One of the challenges of the behaviour-based approach is the integration and co-ordination of multiple different behaviours. This problem was faced using combination operators applied to lower level behaviours and producing useful and coherent higher level behaviours. Our autonomous mobile robot also exhibits learning capabilities, namely the ability to learn new behaviours in an autonomous way. In this context it was designed to the elementary behaviours a connectionist architecture that encompasses the capabilities of building specific basic competences. Through the exploration of the proposed control architecture a mobile platform has become an useful autonomous mobile robot. It is able to deal with dynamic environments, evolving in a coherent, relevant and adequate way according to both the current situation as well as with its assigned missions.

1. Introduction

New potential application areas for mobile platforms have recently emerged, leading to an increasing interest in new architectures for autonomous mobile robots. Moreover, it has been considered a rich field for developing, testing and integration of advanced concepts and approaches based on Artificial Intelligence and Robotics. We should not forget that "the new approaches that have been developed recently for Artificial Intelligence arose out of the work with mobile robots" in (Brooks 1991a).

The first works on intelligent robotics used the

traditional perspective of Artificial Intelligence. In this context the intelligence is viewed as a process of manipulating explicit representations of the world, and the mechanisms for intelligence were algorithmic. This approach, which has been called SMPA (Sense - Model - Plan - Act) approach, has some important drawbacks. It leads to robots operating quite slowly and when placed in dynamic environments, due to some unpredictable situations, their initial plans can be unapplicable.

It was in this context that reactive approaches arose to deal with tasks where the responses must be indexed directly over the situation description rather than resulting from complex problem solving using abstract world models.

In the middle of eighties, the M.I.T. Mobile Robot Group, based on simple biological systems has proposed a new approach - Behaviour-Based Artificial Intelligence. This approach has produced new models of intelligence. The intelligence should provide to the system, situated in an environment, the capability to be reactive to its dynamism and also to generate robust behaviour in the presence of an uncertain, unpredictable and changing real world. The Subsumption Architecture (Brooks 1986) has been described as one that combines the reactivity of agent responses with distribution of the actions' generation.

The Control System architecture presented in this paper embodies key ideas coming from the behaviour-based approach and its purpose is to make it possible to the robot to exhibit the following features:

- capability to react fast and appropriately to the situation always having its goals in mind;
- capability to deal with any possible situation and never blocking - robustness;
- capability of extendibility and dynamic adaptation in order to improve its performance through the use of learning capabilities.

Our main contributions has been concentrated on two areas. One is related to the merging of the solutions proposed by the active behaviours with the orientations proposed by cognitive level agents. The other area focus on the learning of new competencies.

Permission to make digital/hard copies of all or part this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

It seems important to refer that in our approach the concept of agent is present at two different levels. The mobile platform with its specific control system is viewed as an autonomous agent able to deal with the dynamic environment carrying out its mission. In other side the control system is viewed as a society of agents that cooperate to allow the mobile platform to exhibit an appropriate performance.

This paper is organised as follows. After a brief presentation of the system, in section 2, the section 3 presents the Control System architecture as a whole. The section 4 gives a more detailed description of the Behaviours' Community which is the kernel of the proposed controller architecture. The elementary behaviours as well as their integration on higher level behaviours is also presented. In order to promote the possibility of an autonomous acquisition of new competences, we propose, in section 5, a connectionist architecture implementing elementary behaviours which allows the robot to acquire new competences. In section 6 the testbed platform is described, and finally in section 7 some conclusions are drawn and some directions for further research are suggested.

2. The Autonomous Robot Agent Structure

The control system proposed aims to control a mobile robot platform evolving in a dynamic environment behaving as a situated and embodied autonomous agent. Situated because the robot agent deals directly with the world and not just with an abstract description of it. Embodied because the robot has body and its actions take place in the world and have immediate feedback on the agent (Brooks 1991). We claim that our mobile robot may evolve in dynamic environments performing in a coherent, relevant and adequate way according to both the actual situation as well as its assigned mission.

The Autonomous Robot Agent (ARA), depicted in Fig.1, is made up of a Mobile Platform, a Sensors/Actuators Interface and a specific Control System. The Sensors/Actuators Interface is the module that establishes the interface between the Control System and the Mobile Robot Platform. The testbed Mobile Robot Platform is presented in detail in section 6. The Control System (CS) is the most important component of the ARA and its description is the subject of this paper.

The usefulness of an autonomous agent is its ability to perform an assigned mission. In this approach we have a Mediator module whose purpose is to interact with both the user and the ARA. Firstly, by interaction with the user the Mediator gets the descriptions of the mission to be executed. According to the specified mission and through the access to the Knowledge Base the Mediator installs the cognitive and behaviour agents, i.e. the ARA competences,

which will give to the mobile robot the ability to execute the particular mission. Whenever ARA has not enough competences to deal with a particular situation it can ask for helping to the Mediator which search for appropriate competences in the KB, or interact with the user to redefine the mission.

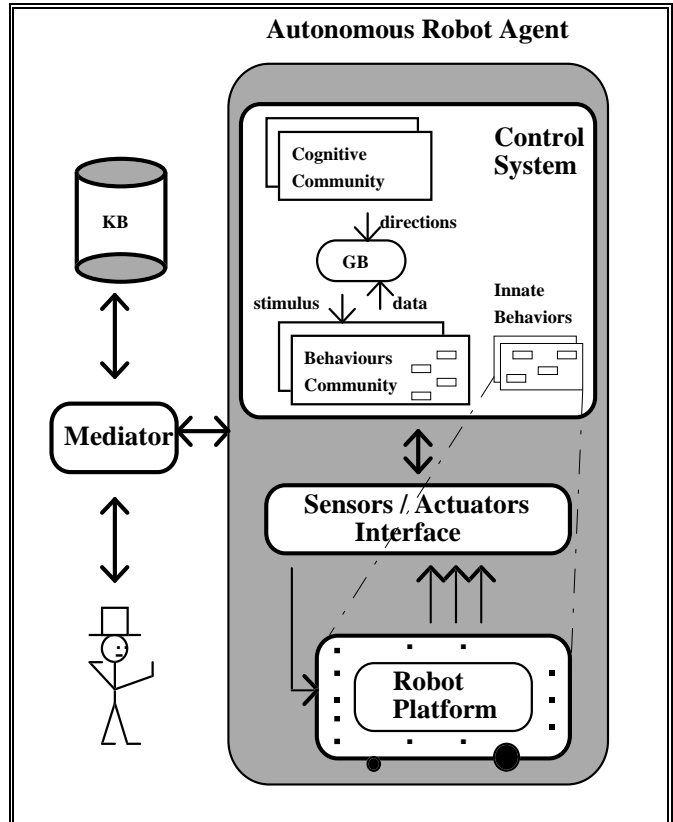


Fig. 1 - The Autonomous Robot Agent Structure

In the Knowledge Base there is knowledge about all the competences of the ARA and also how to install them. Whenever there is topological and geometrical information about the region where the mission will take place, it will also reside in KB, and it will support the cognitive agents in their adviser role.

In a next phase of our project a methodology (Colombetti 1996) to support users in the definition of new missions will be developed.

3. Control System Architecture

The Control System Architecture includes important aspects of reactivity and distribution of the Subsumption Architecture (Brooks 1986), more specifically the key ideas of the Behaviour-Based approach ((Brooks 1991b) and others additional ideas (Maes 1989), (Connell 1990), (Mataric 1994)).

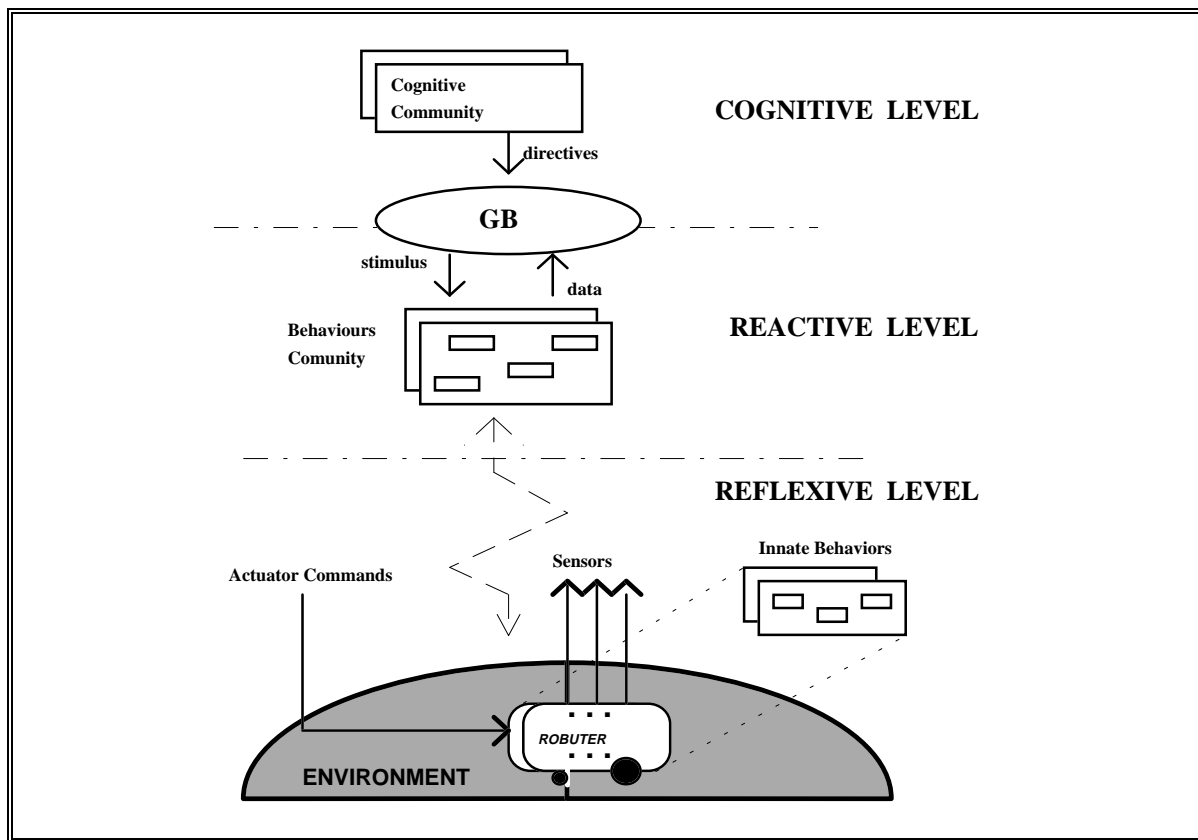


Fig. 2 - Control System Architecture

It was also designed in the scope of the theory "The Society of Mind" (Minsky 1985), where the system is viewed as a society of agents, each of them with its specific competence, that collaborate in order to guide the society in the process of achieving their goals. In this perspective the Control System can be considered a society of competence modules, the agents, that cooperate to allow the autonomous mobile robot to achieve its goals.

We intend, under this modular and distributed Control System to build up a situated and embodied autonomous mobile robot, i.e. able to deal directly with the changing world, which exhibits a certain degree of intelligence essentially emerging from the interactions with the world. Fig.2 shows the organisation of the most relevant agents.

In the proposed distributed architecture the constituent competence modules - the agents - are organised in three levels, the Reflexive Level, the Reactive Level and the Cognitive Level.

The Reflexive Level includes the behaviours responsible for the most elementary competences, the Innate Behaviours. These behaviours are considered reflexive because they act in a pure stimulus-response way (Goodwin 93).

To the Reactive Level belongs the Behaviours' Community implemented as a set of agents which are able to answer quickly to the events in the world because they don't have to deal with heavy computation.

To the Cognitive Level belong the Community of the Cognitive Agents and the Green Board Agent (GB). The principal role played by the cognitive level agents is to advise the Behaviours' Community allowing the ARA to exhibit an "oriented" reaction.

The Green Board Agent acts as an interface agent between the Behaviour's and Cognitive Communities. It has the ability to convert the directives from the cognitive agents into specific stimulus to convey to the behaviours to which they are relevant.

If no topological and geographical information about the place where the mission will take place is available the Cognitive Agents Community will be empty.

In the following sub-sections we will describe more in detail the meaning of each one of the elements in the proposed architecture.

3.1 Innate Behaviours

Innate Behaviours are represented by agents responsible for the most elementary competences also considered survival

instincts. The Innate Behaviours, which are already implemented, include the behaviours "observing", "not-colliding" and "dodging".

The "observing" behaviour corresponds to the sensors reading and gives to the robot the ability to "see" around.

The "not-colliding" behaviour has the competence of collision avoidance. Whenever an obstacle is detected in the movement's direction the robot stops immediately

The "dodging" behaviour corresponds to the instinctive reaction to a quite close and unexpected event. The reaction consists in a light deviation from the obstacle to the opposite side of the direction of the approximation.

These agents belong to the most basic level in order to emphasise the idea that they correspond to basic survival instincts. They reside as close as possible of the robot hardware and should be considered a robot "instinctive" feature. This is the reason why they are called Innate Behaviours. These set of agents are always present in all the missions. They are always active having higher priority, which means that their associated actions are always immediately executed.

3.2 Behaviours' Community

The Behaviours' Community is the cornerstone of the proposed Control System architecture. Based on the Subsumption Architecture and extended with additional features that aim to improve the system's performance, it is responsible for the claimed capabilities of reactivity, distributivity, robustness, extendibility and adaptability of the autonomous mobile robot. The agents of this community compete, at any moment, for controlling the robot.

Each behaviour of this community is responsible for a specific competence and it can be considered an agent because, together the Innate Behaviours, it controls the mobile robot in carrying out this specific competence. It perceives its environment, through the sensors' information, and decides about the action to be done converting this decisions into appropriate commands to the actuators.

Due to the Behaviours' Community importance, further details will be given in section 4.

3.3 Cognitive Agents' Community

Considering that knowledge, used in an appropriate way, can improve the execution performance of a robot, without reducing the capabilities to react to the environment, we believe that cognitive agents could have an important meta-level role in advising for a better reaction.

Cognitive agents, accessing the topological and geographical information in the Knowledge Base and according to the current mission, can develop alternative abstract plans without going into details. Building several

alternative abstract plans prevents the robot from losing time in replanning.

These abstract plans will be detailed during execution time enabling the robot to display "directive reaction" instead of getting "directed control from a plan" (Budenske 92).

The cognitive agents should also be able to collect the relevant information, present in the developing of abstract plans in order to provide it to the agents' society. In this way the loss of information that abstraction can cause (Payton, Rosenblatt, Keirse 95) is avoided, and the robot is able to take opportunistic advantages of unexpected situations. This information is used to build up, associated to each abstract plans' step, the correspondent artificial potential field where the goal locations exert an attractive force and the fixed obstacles exert a repulsive one. From these artificial potential fields pseudo-gradient field maps, which attract the robot to the goal locations, are developed. These maps lead to specific directives which are sent to the Green Board Agent (GB) in order to be transformed into advises to the Behaviours' Community. The oscillations around a position that can occur when a robot, guided by these kind of maps, finds out a new obstacle, are avoided by the GB because it never advises the robot to go to recent position.

Between the Cognitive Level Agents and Behaviours' Community a kind of co-operation which we call advising co-operation is established.

It can also be interesting to add another kind of cognitive agents able to build up and continuously update maps like: plants, occupied space probabilistic maps and so on.

3.4 Green Board Agent

This agent is a special kind of blackboard. From the mediator it receives the description of the mission and from the cognitive agents, when knowledge about the mission is available, it gets the alternatives abstract plans plus the associated pseudo-gradient field maps. All this information becomes available on the board.

According to the current situation characterised by sensors' information, by the actions already performed and having in mind the goals, GB processes these information and provides it as supplementary stimulus to the behaviours. These stimulus are sent to the active behaviours to reinforce the selection process of the potential actions.

The GB can also be considered an interface agent that collects the directives issued by the cognitive agents, transforms them in stimulus, according to the current situation. These stimulus are provided to the behaviours to which they are relevant. These stimulus are used as non-mandatory information to advise the reactive level rather than as strong constraints imposed upon reactive level by the cognitive one.

4. Behaviours' Community

As it was referred in section 3, the Behaviours' Community is the kernel of the proposed autonomous mobile robot Control System. The agents of this community together with the Innate Behaviours are enough to guarantee a flexible control of the robot. The cognitive level agents are responsible for enhancing the performance of this agents' community, playing a meta-level role as adviser, providing information that will act as supplementary stimulus to guide the robot actions.

Two kind of agents inhabit in this community, the basic or elementary behaviours and the compound or high-level behaviours. The basic behaviours perform tasks considered elementary and the high-level behaviours are responsible for more complex tasks involving the integration and co-ordination of more elementary ones. In this perspective the basic behaviours will be the substratum for a variety of high-level behaviours which by means of different combinations of those ones produce the desired composite output behaviour (Mataric 94).

The generic architecture of each behaviour, inspired by the work of (Maes89), (Connell 90) and (Correia 95), is essentially composed by two modules. The Activation Module determines the agent activation level which is an important value to decide about the agent activation when it is integrated in a higher level behaviour. The Action Module defines what sort of action may be taken based on the current input.

The input of each behaviour can include data from the sensors and other potential stimulus coming from the Green Board agent or from the behaviour that includes it.

There are behaviours that, when are faced with a situation, can find several alternative solutions, some probably more attractive than others. To each potential solution is possible to associate an adequacy level that allows to score them attending to their convenience for the situation in the behaviour perspective. In order to decrease the possible hidden information in the behaviours a multiple output was chosen. This multiple output is a set of couples (adequacy - action) that become accessible, when necessary, all the available hypothesis.

When a behaviour is integrated in a higher level behaviour its multiple output flows to the higher behaviour and from the potential actions the effective action is selected.

4.1 The Elementary Behaviours

Elementary Behaviours are defined in accordance to some particular competences we want to achieve. They must be simple, correct, robust and moreover each of them must act as an independent and self-sufficient agent in performing its own specific task.

The selection process of the elementary behaviours set is not easy and it is constrained by the robot platform dynamics, by the environment and also by the robot's goals specified at the mission level.

This selected set of behaviours should be small, but enough to enable the execution of complex tasks.

In this phase all the behaviours were designed by programming the strategy that maps perception to action. In sub-section 5.1 another approach is presented which allows autonomous acquisition of competences.

4.2 Integrating and Co-ordinating different behaviours in High-Level Behaviours

The problem of co-ordinating the activity of different behaviours is one of the challenges of behaviour-based Control Systems and different approaches solve it in different ways.

In this architecture lower behaviours are combined by specific combination operators to generate more complex ones, called high-level behaviours. Each one of the operators represent a collaboration strategy among the argument behaviours. Four operators were defined but others can be considered due to new needs and they can easily be included in this same architecture.

The "+" operator specifies a combination of mutually exclusive behaviours. Due to mutually exclusive behaviours are activated by mutually exclusive conditions, when a behaviour is activated it takes the full robot control.




Fig.4. - Behaviours' combination trough "+" operator

The operator ">" specifies a weighed combination of behaviours having as result the output proposed by the higher activated behaviour.




Fig.5. - Behaviours' combination trough ">" operator

Each behaviour has its activation function, which combines information from sensors, information about the current state and other factors, and determining the activation level of that particular behaviour. Fatigue factors are also sometimes associated to some behaviours in order to avoid impasse situations. The value of activation function called activation level is considered the factor that scores the behaviours. When a high-level behaviour is defined through this operator, a priority factor must be assigned to each one of the component behaviours.

The operator "&" specifies a combination of multiple concurrently active behaviours. The output satisfies all the component behaviours because it is the intersection of every proposed outputs.




Fig.6. - Behaviours' combination through "&" operator

The operator "#" specifies a combination of behaviours characterised by activation of one of the component behaviour. According to a set of trigger conditions (Cond.) just one of the behaviours can be activated at a time. Trigger conditions are predicates on data that characterise the current situation. The strategy implemented by this operator is a switching over behaviours assuring a coherent sequence of outputs.




Fig.7. - Behaviours' combination through "#" operator

Following this approach a new programming paradigm can arise to define a composite behaviour.

4.3 A Brief Example

Suppose that were defined the elementary behaviours "moving", "attracting-to", "aligning-to", "passage-detector" and "crossing". They are considered as sufficient and

necessary to achieve the potential assigned missions. The "moving" behaviour corresponds to the capability of moving while avoiding obstacles. The "attracting" behaviour implements the ability to go straight a position which works like an attractive point. The "aligning-to" is related with the competence to align with a static obstacle like a wall. The "passage-detector" is continuously paying attention to possible passages during the robot movement. And "crossing" is the competence to cross a passage. From this set of elementary behaviours were designed high-level behaviours.

The behaviour "going-to" which gives to the robot the ability to achieve a defined position, was built by combining the behaviours "moving" and "attracting" through the "&" operator. The action of this higher-level behaviour is the one that better satisfies both behaviours.

$$\text{going-to} = \& \{ \text{moving}, \text{attracting-to} \}$$

The behaviour "following" was built in a similar way

$$\text{following} = \& \{ \text{moving}, \text{aligning-to} \}$$

The behaviour "finding-passage" can be considered a ">" combination of "passage-detector" and "following". A greater priority should be assigned to the "passage-detector". When a passage is found the true value is assigned to the global variable "v_passage" and the action proposed by "passage-detector" is stop.

$$\text{finding-passage} = > \{ (5, \text{following}), (10, \text{passage-detector}) \}$$

The behaviour "cross-passage" can be built by applying the "#" combination operator to the behaviours "finding-passage" and "crossing".

$$\text{cross-passage} = \# \{ \text{finding-passage}, \text{crossing} \}$$

The associated trigger conditions and behaviours table is showed below.

CONDITIONS	BEHAVIOUR
v_passage	
0	finding-passage
1	crossing

5. Learning Capabilities

Learning capabilities has the propose to facilitate the possible actions of the autonomous agent making them more relevant, appropriate and precise.

Learning of new competences, and learning to co-ordinate existing behaviours are two learning capabilities we are trying to include in the control architecture of our robot for the sake of the increase of its autonomy.

5.1 Autonomous Acquisition of New Elementary Competences

There are at least two different ways to make the autonomous mobile robot to execute a specified task. One way is to pre-program the strategy that maps perception into action. The other way is to allow the autonomous mobile robot to acquire its competences in an autonomous way through a learning capability. The last alternative is the most appropriate when it is difficult, bored or even impossible to apriori specify the perception action strategy.

The architecture proposed in this section, depicted in Fig.8, aims to achieve this last goal.

This architecture is based on a simple connectionist paradigm where perceptions are associated to actions by means of an artificial neural network (Nehmzow 1992). A fundamental characteristic of this approach is the reduced use of external knowledge.

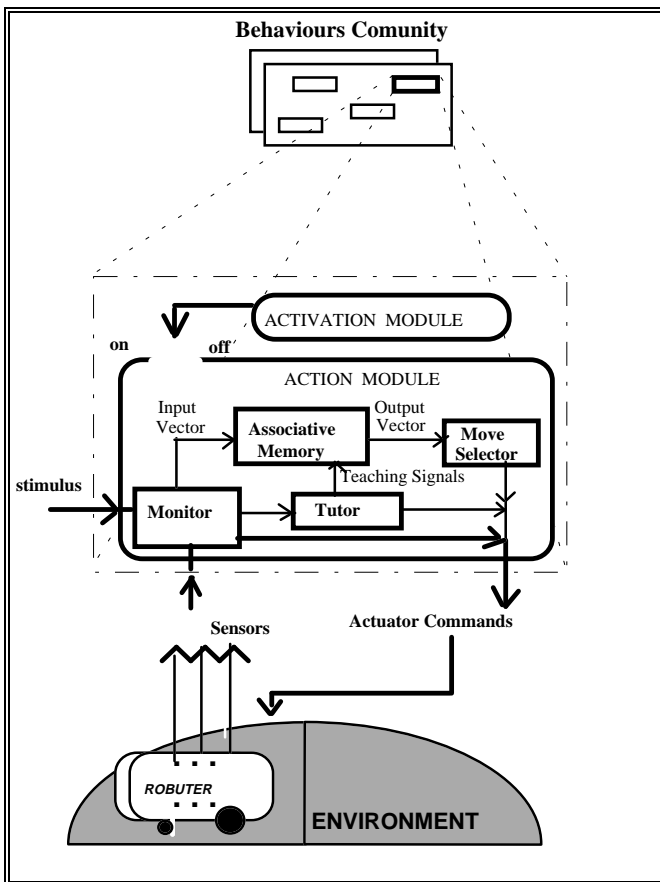


Fig.8 -Elementary Behaviour Architecture

The central component of this behaviour agent architecture is an associative memory implemented by an artificial neural network. Inputs to the this neural network are pre-processed perceptions that can include sensor values and other variables. Output signals of the associative

memory represent the actions corresponding to those input perceptions.

In order to achieve a very fast learning, the neural network should be as simple as possible. In robotics there are a lot of linearly separable functions where a perceptron, the simpler neural network, can learn successfully the mapping function between inputs and outputs.

The monitor is the module where the external knowledge resides. This knowledge is represented by a small set of rules that expresses believes and desires of the behaviour agent (Wooldridge, Jennings 95). This behaviour agent is always working according to its believes and desires. When a "tension situation" - situation in which there is no knowledge to deal with - occurs, an input vector to the neural network is generated and an output is derived in order to overcome this tension situation. If the neural network does not have any answer to the situation, the tutor module generates a possible action or a set of possible actions. If these generated actions contribute to, immediately or in a short term, overcome this situation then they will be learned by the neural network. A "tension situation" is considered overcome by an action or a sequence of actions if these actions led the robot to a new situation to which it has knowledge to deal with. Therefore, from the "tension situation" and the correspondent actions, teaching signals are produced and lead to the adaptation of the network.

Depending on the behaviour that is being learned and the environment conditions, it can be useful to substitute the tutor module by a human tutor that advises the action or actions' sequence to overcome the tension situation. In this way the learning process becomes faster but the complete autonomy of the agent is lost.

Trough trial and error strategy the behaviour agent will be able to associate perceptions to actions, building up an appropriate associative memory. Ideally, the learning phase should be over when the robot has met every possible situation at least once, but preferably several times. In practice if the robot has not made a mistake in the choice of the action during an enough number of times, we consider the competence as learned.

When the competence is considered as learned, the acquired knowledge in the associative network is frozen and the agent architecture is simplified and, therefore becoming more efficient.

6. The Testbed Platform

This Control System has been tested in the mobile platform ROBUTER II (Robotsoft SA). This is a rectangular platform with 1025 mm length, 680 mm width, 440 mm height and about 150 Kg. It has two rear independent propulsive wheels with differential steering and two front free wheels. Onboard, it has a computer compatible with

VME Bus, based on microprocessor 68020, 32 bits, 16 MHz, 1 Mbyte of RAM, running the real time operating system ALBATROS, specially designed for multi-axis control.

It allows interactive communication from a host machine, by sending and receiving ASCII messages through a wireless modem. It also allows the developing and downloading of an application program, codified in C language, from the host machine. The host machine used is a Sun SPARC Classic with SunOS 4.1.3.

The platform is equipped with a complete hardware and software ultrasonic system that includes a dedicated CPU controlling those ultrasonic sensors. There are 12 ultrasonic sensors (5 placed in front, 2 on each side and 3 on back). Two capacitive proximity switches are also installed in front on each side.

It also has an odometric system providing an estimation of the robot position in a Cartesian plan.

We have already implemented and tested, with this platform the innate behaviours, being the behaviour community under development.

The innate behaviours were implemented as resident processes in the robot on-board computer running periodically and managed by the robot operating system.

The behaviour community is running in the Sun station and interact with the robot via radio (Fig. 2). This option was made after several tests which prove that the robot performance decrease rapidly when many processes are running simultaneously in the on-board computer.

7. Conclusion

With the described architecture our mobile robot platform can now be seen as a real autonomous agent with the following features: situation and goal orientedness, robustness and extendibility.

Situation and goal orientedness, also referred as oriented reaction, is related with the robot capability to react to the environment changes directed by its goals. If pre-existing knowledge about the mission is available it will be used for advising.

The robustness capability means to deal with any situation without blocking. This capability is achieved because there is always at least one behaviour which is able to take the robot control even something fails.

Due to the behavioural approach of the system, it is easy to include more competences. The formalization of the combination operators in order to provide the user a language for defining new competences is the next step.

Learning capabilities allowing the robot to adapt itself and better deal with new and different situations has also been one of our concerns.

References

- (Brooks 1986) - Rodney A. Brooks - " A Robusted Layered Control System for a Mobile Robot " in IEE Journal of Robotics and Automation Vol RA2, nº1 March 1986
- (Brooks 1991a) - Rodney A. Brooks - " Intelligence Without Reason" in IJCAI-91 - 1991
- (Brooks 1991b) - Rodney A. Brooks - "Integrated Systems Based on Behaviors" in SIGART Bulletin Vol.2, No. 4 - 1991
- (Brooks 1991b) - Rodney A. Brooks - "New Approaches to Robotics" in Science, 253:1227-1232, 1991
- (Budenske 1992) - John Budenske - " Knowledge, Execution and What Sufficiently Describes Sensors and Actuators" in Mobile Robots VII - 1992
- (Colombetti 1996) - Marco Colombetti, Marco Dorigo, Giuseppe Borghi - "Behavior Analysis and Training - A Methodology for Behavior Engineering" in IEEE Transactions on Systems, Man, and Cybernetics, Vol.26. NO.6,1996
- (Connell 1990) - Jonathan Connell - "Minimalist Mobile Robotics - A Colony-style Architecture for an Artificial Creature" Academic Press - 1990
- (Correia 1995) - Luís M.P.Correia - " Veículos Autónomos Baseados em Comportamentos" - PhD thesis - UNL 1995
- (Goodwin 1993) - Richard Goodwin - "Formalizing Properties of Agents" - May 93 CMU-CS-93-159
- (Maes 1989) - Pattie Maes - "How To Do the Right Thing" A.I. Memo 1180 - M.I.T,1989
- (Mataric 1994) - Maja J. Mataric - " Intercation and Intelligent Behavior " - Ph. D. Thesis- M.I.T. , May 1994
- (Minsky 1985) - Marvin Minsky - "The Society of Mind" Simon & Schuster, 1985
- (Nehmzow 1992) - Ulrich Nehmzow - "Experiments in Competence Acquisition for Autonomous Mobile Robots" Ph.D. - University of Edinburg 1992
- (Payton, Rosenblatt, Keirse 1995) - David Payton, Kenneth Rosenblatt, David Keirse - "Plan Guided Reaction" in Autonomous Mobile Robots- Control, Planning and Architecture Vol2 edited by Iyengar and Elfes Nov 88
- (Robotsoft SA)- "ROBUTER Users Manual - Robotsoft SA - April 1993
- (Wooldridge, Jennings 95) - Michael Wooldridge, Nicholas R. Jennings - "Agent Theories, Architectures, and Languages: A Survey" in Wooldriges and Jennings Eds., Intelligent Agents, Berlin : Springer-Verlag, 1-22.