RoMan '97. Sendai, October 1997

A Human Interface for an Intelligent Mobile Robot

Volker Graefe and Rainer Bischoff

Institut für Meßtechnik
Universität der Bundeswehr München
85577 Neubiberg, Germany
Fax: +49 89 6004-3074, e-mail: Graefe@UniBw-Muenchen.de

Abstract

We propose that basing robot control on situation assessment may lead to the realization of a truly user-friendly human-robot interface. Two aspects of user friendliness are addressed: ease of introducing domain-specific knowledge into the robot's memory in preparation of its actual operation, and ease of giving orders to the robot.

Two abilities of the robot that result from the situation-oriented approach are key to this user friendliness: its ability to acquire knowledge regarding the environment largely by learning, instead of depending on explicit inputs by an operator, and its ability to communicate with the operator on the common basis of the situation being perceived by both. These abilities, together with a specific structure of the robot's internal knowledge representation, allow the operator to communicate with the robot in a similar way as he would with a human whom he requests, for instance, to do some errand.

1. Introduction

Future service robots will have to interact closely with humans and, specifically, with humans who are not robotics experts and often not even interested in robotics and other technical matters. Moreover, if such robots are to be deployed in massive numbers, e.g., in unstructured and continuously changing household environments or construction sites, it is clearly impossible to have professional experts spend much time for setting up each individual robot and adjusting it to the characteristics of its particular environment.

A high degree of user friendliness and, specifically, a human interface that allows the user to communicate with the robot in a pleasant and intuitive way, as well as the ability to operate in unforeseen and variable environments without depending on any individual training by an expert are, therefore, indispensable requirements for future personal robots. In the sequel we will discuss approaches that we expect to contribute to developments in this direction and report some results that we have achieved with them in experiments with a mobile robot in an unmodified laboratory environment. We will concentrate on the discussion of two characteristics of user-friendly service robots: the ability of acquiring domain-specific knowledge by learning instead of needing programming or training by an expert, and the ability to accept orders of a similar form and abstraction level as orders that a human would give to another human or to a domestic animal.

We will show that a certain behavior-based robot architecture that relies on an understanding of situations for the selection of behaviors to be executed is particularly suitable for achieving such characteristics. A robot implementing this architecture allows a human interface to be realized that makes the robot appear intelligent and easy to communicate with.

After introducing the concept of "situation" and a robot architecture incorporating this concept we will describe a vision-guided autonomous mobile robot that we have built according to this architecture. The robot is able to build a map of the environment largely automatically by supervised learning. Using this map it then navigates in the charted environment and executes orders given on a relatively high abstraction level, such as "go to the electronics lab", referring to place names as they are ordinarily used by humans, rather than to any coordinates or the names of robot-internal variables. This is a direct consequence of its ability to understand situations, in connection with the ways how domain-specific and task-specific knowledge is represented internally.

2. Situation-Oriented Behavior-Based Robot Architecture

Behavior-based system architectures are now generally accepted as an efficient basis for autonomous mobile robots. Their main principle is the achievement of desired goals by activating an appropriate sequence of behaviors. This method allows, for instance, a mobile robot to navigate intelligently in a complex network of passageways without requiring a global coordinate system or an accurate map [Graefe, Wershofen 1991].

A behavior-based robot replaces the time-consuming path planning that is typical for coordinate-based robots by a selection of behaviors from a repertoire of pre-defined behaviors. The key problem in designing behavior-based robots is the question how to choose at each moment the most appropriate behavior. We propose to base the decision on a multitude of factors that we summarize under the term “situation”.

2.1 Situation

Conventionally, the term “situation” describes, among others, “the way in which something is placed in relation to its surroundings”, a “state”, a “relative position or combination of circumstances at a given moment” or “the sum of total internal and external stimuli that act upon an organism within a given time interval” [Babcock 1976]. In the context of a behavior-based robot we define the term “situation” in a more operational way as the set of all decisive factors that should ideally be considered by the robot in selecting the right behavior pattern at a given moment.

These decisive factors are:

- perceivable objects in the environment of the robot and their suspected or recognized states
- the static characteristics of the environment (e.g., as stored in a map), even if they cannot be perceived by the robot’s sensors at the given moment
- the state of the robot (state of motion, presently executed behavior pattern, ...)
- the repertoire of available behaviors and the abilities of the robot to change the present situation in a desired way by executing appropriate behaviors
- the goals of the robot, i.e. permanent goals (survival, obstacle avoidance) and transient goals emerging from the actual mission description (destination, corridor to be used, ...) or directly imposed by the human operator.

If all the factors that constitute a situation are, indeed, taken into account by a robot in its behavior selection to the largest extent that is practically possible we call the architecture of the robot “situation-oriented behavior-based”. We are convinced that basing the selection of behaviors on the perceived situation bears some similarity to the ways how humans decide on their actions and that it has the potential of leading to the design of robots that people can easily and intuitively interact with.

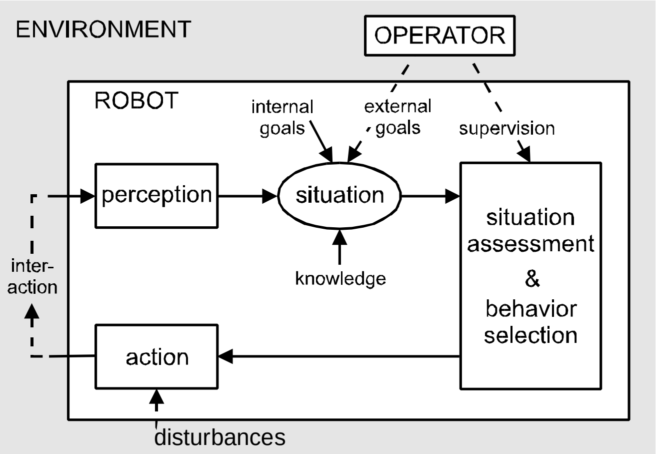

Figure 1:
The role of “situation” as a key concept in the perception-action loop of a situation-oriented behavior-based robot

Figure 1 illustrates how the concept of “situation” may be embedded in the perception-action loop of a behavior-based robot. The figure also makes it clear that the “situation” on which the robot bases its behavior selection is, strictly speaking, only the robot’s internal image of the actual situation. Due to imperfect sensing or imperfect knowledge, this image may sometimes differ from the true situation, which will then result in a sub-optimal, or even grossly inappropriate, behavior of the robot.

After assessing the situation a behavior is selected and executed. Actually, the selection of a behavior that is appropriate for the situation may be considered the main purpose of the situation recognition. The actions of the robot normally change the state of the environment (or world), and some of these changes are perceived by the robot’s sensors. This closes the perception-action loop. Disturbances during the behavior execution lead to non-expected situations and can be corrected by either adjusting behavior-immanent parameters or selecting a different behavior.

2.2 System Architecture

Figure 2 shows the essence of the situation-oriented robot architecture as we have implemented it [Wershofen 1996]. The most important module, the situation module, interacts with the sensors, the actuators, the data base management and the man machine interface in a bidirectional way. Fusing data and information from the different modules makes situation assessment and behavior selection possible. Acting as the core of the whole system this module is also responsible for system management and behavior coordination, including the coordination of all communication over the man-machine interface.

As in all behavior-based architectures, the actions of the robot are based on a repertoire of built-in behaviors. The execution of each behavior is supervised by a specific coordination process within the situation module. By activating and deactivating these coordination processes a management process realizes the situation-dependent concatenation of elementary behavior patterns to complex and elaborate action.

2.3 Knowledge Representation

Since the way how knowledge is represented in a robot is crucial for the characteristics of the man-machine communication, we describe the knowledge representation that meets the requirements of the situation-oriented behavior-based architecture in a little more detail. We will show that this type of knowledge representation is also an excellent basis for a user-friendly human interface.

Two different types of knowledge are needed by an autonomous mobile robot: knowledge of the mission to accomplish and knowledge of the static characteristics of the environment (geographical knowledge). Accordingly, as Figure 2 shows, two knowledge bases exist in the robot, a mission description and a map.

Each knowledge base consists of two parts, stored data and a so-called "expert". Maintenance and management of the stored data and retrieving bits of knowledge at the situation module's requests are the tasks of the "experts".

The attributed topological map

An attributed topological map has proven to be particularly suitable for storing the geographical knowledge needed by a behavior-based mobile robot [Graefe, Wershofen 1991]. Figure 3 illustrates the concept by showing a small section of such a map. Passageways are modeled as pairs of directed paths, while intersections and junctions (A, B, C ...) of passageways, and task relevant locations, such as docking stations (e.g. G), are represented as named points.

Each passageway connects two of these points. Depending on the direction in which a passageway is traversed, different objects will be visible and relevant for the robot, and different maneuvers will be indicated. Therefore, two directed paths, corresponding to the two possible directions of traversal, are assigned to each passageway, each having its own attribute list. For example, in Figure 3 the two lists, AF and FA, are assigned to the two directed paths connecting the points, A and F.

Typical attributes of a directed path include its approximate length and a list of landmarks or other significant objects presumably visible while traversing the directed path. For each of these objects information is included as to its approximate location relative to the passageway, how it may be recognized in an image, and possibly its name. The typical attributes of a named point include lists of all directly adjoining named points and of the approximate angles between the passageways leading into the point.

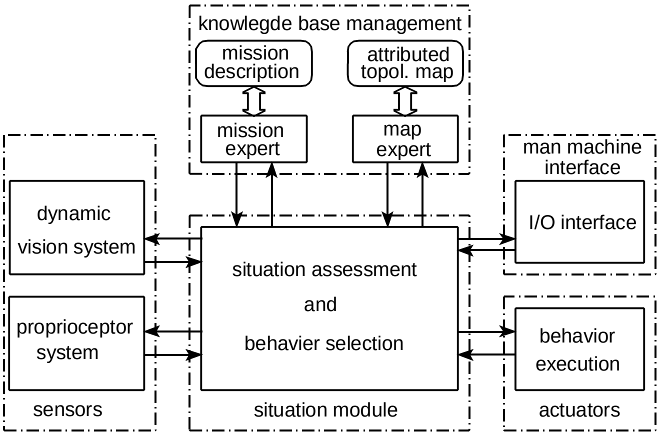

Figure 2:
System architecture of our situation-oriented behavior-based mobile robot

Also indicated in the map may be certain behaviors or maneuvers that should always be executed when a particular path or location is being traversed. Actually, there is almost no limit to the variety of information that may be included in such a map as attributes. This flexibility is a great advantage of the attributed topological map.

The map must represent the topology of the passageway network correctly, but as for the geometrical attributes, it suffices if all angles and distances are only approximately correct. This semi-qualitative and coordinate-free nature is another great advantage of the attributed topological map. It makes this type of map very suitable for the behavior-based approach; moreover, due to its nature, it is much easier to generate than a geometrically accurate map as required by

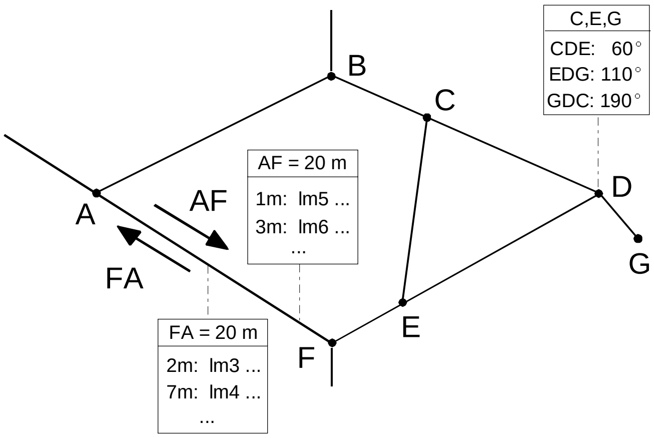

Figure 3
Section of an attributed topological map. Attributes (here shown in an abbreviated form) with hints on geometrical quantities, visible landmarks, lm, etc. may be attached to each directed path and each named point.

some other navigation methods (e.g. coordinate-based navigation, cf. [Graefe, Wershofen 1991]). In fact, because it is so easy to generate, a robot may generate such a map by learning with a minimum of operator assistance (cf. section 4).

The mission description

Mission descriptions for a situation-oriented behavior-based robot may have many forms. On the lowest level a mission description is a list of built-in behaviors with all necessary parameters. A mission specified in this form may be performed by the robot by simply executing the listed behaviors, where a successful termination of one behavior triggers the execution of the next one in the list. While this type of low-level mission description is suitable for immediate execution by the robot, a human operator normally prefers to specify the mission in a much more compact form and on a higher level of abstraction.

From the user's point of view, in most cases a mission description should ideally consist of only a single command, e.g., "Move to office 1, pick up object O and transport it to office 2!". Depending on the task, the preferences of the user, and his confidence in the planning ability of the robot, many intermediate levels may be desirable, too. For example, to make route planning easier for the robot, a list of intermediate points could be specified, in addition to the final destination.

Obviously, this type of communication with a robot on the basis of high-level task descriptions resembles instructions that might also be given to a person. However, only a mission description that refers directly to the elementary built-in behaviors of the robot can eventually be executed. If a mission description consists originally of other items than these elementary behaviors the robot's mission expert, in cooperation with the map expert, must convert it into a sequence of built-in behaviors which may then be executed. An example will be shown in section 4.

2.4 Human Interface

As Figure 1 shows, the situation-oriented behavior-based architecture provides the operator of the robot with two ways of influencing its behavior. He may impose goals upon the robot ("external goals"), and thus modify the situation, by issuing high-level commands that may define a desired final state without specifying how this state should be reached. An example (implying a still fictitious robot) would be the object- and action-oriented command "Bring me a cup of tea!". Alternatively the operator may, if he so desires, control the behavior selection directly on a much lower level in a more or less detailed way. Examples are directing commands, e.g., "Follow the present corridor until the next corner on the right side!" or "Turn left by 45 degrees!" and intervening commands like "Stop!". Such commands are particularly useful during supervised learning, e.g., when the robot is learning the static characteristics of the environment.

Communication with a robot will be easy, pleasant and efficient for a human if he can communicate with the robot in a similar way as with a human. The knowledge bases described above are designed to support certain aspects of such a communication. Representing "geographic" knowledge in an attributed topological map has two consequences: it makes it practical (as will be shown) for a robot, even with limited sensory abilities, to build its map by learning, thus relieving a human operator from the chores of surveying the environment and inputting masses of data; and it allows the operator to use terms that are natural for a human, such as place names or "second corner", instead of coordinates, when communicating with the robot.

The mission expert, together with the map expert and the situation module, enables the robot to understand orders that might also be given in a similar form to a human. A robot with such communicative abilities allows the user to express his wishes in a variety of ways according to his own preferences and without having to know anything about the robot's internal structure. The user has a wide choice as to the level of detail and the degree of vagueness he wants to employ in his communication with the robot.

3. An Implementation: *ATHENE II*

The proposed concepts have been implemented in a mobile robot, *ATHENE II* (Figure 4). Extensive experiments related to learning and vision-based autonomous navigation were performed with it [Wershofen 96]. *ATHENE II* is a three-wheeled vehicle, about 1.35 m long and 0.72 m wide, propelled and steered by the front wheel at a maximum speed of 1.5 m/s.

Figure 4
The mobile Robot *ATHENE II*

Sensors

The main sensor is a video camera with a lens allowing a field of view of about 45°. This camera is mounted on a one-axis platform able to turn around its vertical (pan-) axis with a range of ±150 ° at a speed of up to 360 °/s.

The recognition and assessment of situations requires rich information on the environment and the internal state of the robot. A large part of this information is provided by the sensor subsystem that consists mainly of a dynamic vision system and odometers in the robot's wheels. The reason for choosing vision as the main sensor modality is twofold: vision is arguably the most powerful and flexible sensor modality in most environments where service robots may be expected to operate [Graefe 1992], and vision is the sensing modality that is primarily used by humans for most locomotion and manipulation tasks. Having the robot use the same sensor modality as the human operator is a good basis for an intuitive and effective man-machine communication.

Robot control

The robot control program runs on a PC and is realized as a finite state machine (Figure 5). After powering up the robot finds itself in the state "Waiting for next mission description". Two different forms of mission descriptions are possible: a name list indicating the final destination and optionally a number of intermediate points, or a list of behaviors that are to be executed sequentially. In case of a name list, a planning module relying on the knowledge bases introduced in section 2.3 converts the given route into a sequence of behaviors (state "Planning"). Planning would also be necessary if during the state "Behavior Execution" an unexpected situation were encountered that required planning a new route, and thus, a new behavior sequence. This feature has not been implemented yet (dotted line in Figure 5), but it would be necessary in situations where, e.g., another task gets a higher priority or obstacles, e.g., closed doors, block the robot's path.

In order to plan a route according to its map, the robot needs to know its present location and orientation. If the operator does not provide this information the robot will acquire it automatically (state "Re-Orientation"). Two methods have been implemented, one based on landmark recognition, the other one based on map matching. Both have been tested successfully, the first one in real-world experiments [Wershofen 1996], the second one in simulations. When the robot has determined its location it starts route planning. While the list of behaviors is being executed (state "Behavior Execution") the human operator can intervene and modify goals or behavior execution if he so desires (state "Interpretation of Operator Intervention"). When a mission is completed the robot halts and waits for new orders.

4. Results

Two types of experiments that we have conducted with the robot *ATHENE II* may serve to demonstrate the user-friendliness of the robot's human interface that results from the underlying system architecture. They relate to the acquisition of geographic knowledge by supervised learning with only a minimum of operator effort and to the performance of a complex navigation task by the robot after it had received a relatively simple order.

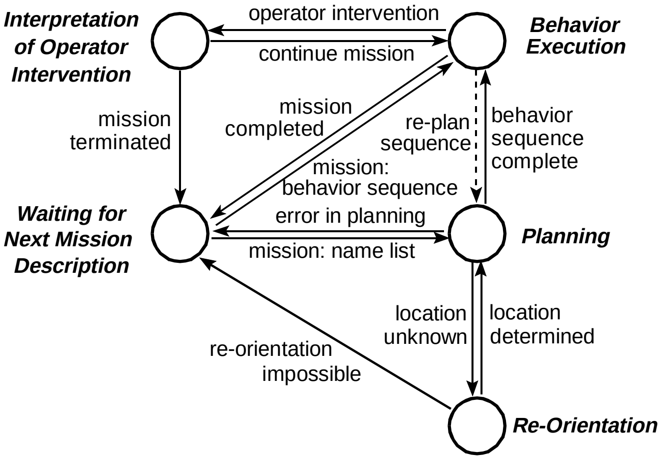

Figure 5
Realization of the robot control software as a finite state machine

Learning

For all service tasks (e.g. mail delivery or cleaning) domain specific knowledge must be introduced into the robot. This training or adaption phase is inevitable if the robot should have a purpose other than research (cf. subsumption architecture imitating insect level behavior without any application purposes [Brooks 1996]).There are two alternatives for introducing domain specific data into the robot: First, knowledge could be entered by experts who program the robot in advance, e.g., based on data from an accurate map. Second, the robot collects and interprets the necessary data autonomously, like a human would do who prepares for the same task. A mail delivery robot could, e.g., wander around and gather information from door plates and signs in the beginning. After a while it would be able to do first deliveries, thereby updating and extending its knowledge base.

Although it would be desirable to have such a robot, the sensory abilities of our robot are still insufficient to master this task if set in an unknown environment. However, it is the advantage of our situation-oriented behavior-based approach that the robot may fairly easily be trained to work in previously unknown environments by a teacher who has some knowledge about the robot's future working environment and a basic understanding of its repertoire of behaviors.

The goal in the experiments was to have the robot generate, or extend, an attributed topological map of the environment. During supervised learning the robot explores an unknown environment together with a human teacher. The geometric information is provided by the robot's dead-reckoning system, and relevant location names, as well as attributes relating to landmarks, are entered by the teacher. When the cognitive abilities of the robot are not sufficient to navigate safely at critical points the teacher provides the correct behavior or action pattern to solve this navigational problem. In this way maps of networks of corridors could be generated that could later be used by the robot for navigating. The time required for the supervised learning was insignificantly longer than the time needed for merely driving the robot through the same corridors by manual control.

Executing a complex navigation task

Figure 6 shows, as an example, the mission description that the robot was given in an experiment, and the resulting course followed by the robot. To make the task more complex for demonstration purposes the robot was instructed to pass a rather large number of intermediate locations on its way to its final destination, the e-lab. The mission description is simply a list of all the locations that should be passed by the robot, and it ends with the final destination.

At the start of the experiment the robot knew that it was somewhere between the e-lab and the kitchen, facing the kitchen. It had a map of the environment that it had acquired in previous experiments, and that did not contain any gross errors in its metric attributes. (In other experiments the robot completed similar missions with maps into which errors of several meters for the lengths of some corridors had been introduced.)

5. Conclusions and Outlook

Communication between a human and a robot should be intuitive, pleasant and efficient for the human. This requires a certain degree of intelligence and user friendliness of the robot. A situation-oriented behavior-based architecture has been introduced that allows robots with such desirable properties to be designed. As other behavior-based robots they perform complex tasks by concatenating elementary built-in behaviors, but in contrast to previous approaches they base the selection of their behaviors on the situation they perceive at that moment. Obviously, this concept requires the robot to be able to perceive a situation and to assess it in real time, which places high demands on its sensing and information processing. Primarily a vision system is able to provide the rich information necessary for a meaningful situation assessment.

Experiments with a vision-based mobile robot have confirmed the validity of the concept. Due to the innate characteristics of the situation-oriented behavior-based approach the robot is able to cooperate with a human in the acquisition of geographical knowledge and to accept orders that would be given to a human in a similar way. Just as a human, the robot does not expect its orders to be expressed in terms of coordinates or exact distances; place names and semi-qualitative references, such as “at the next intersection” suffice.

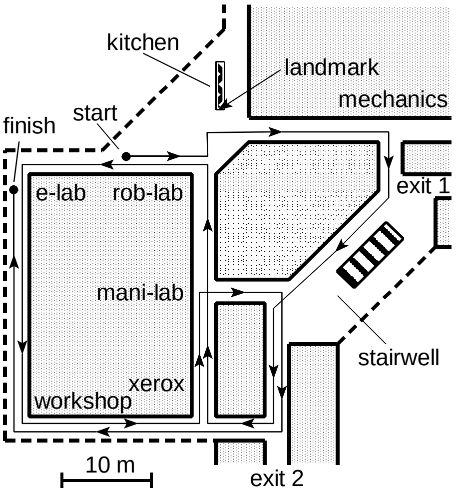

Figure 6
The course traveled by *ATHENE II* according to the mission description shown on the right

NAME LIST
>>>Tour<<<
***
mechanics
exit 1
exit 2
mani-lab
rob-lab
e-lab
mani-lab
stairwell
exit 2
workshop
e-lab
***
END
***

A new, humanoid robot based on the same concepts is currently being built. It has 18 degrees of freedom and is in many respects much more complex than the present one. In view of this complexity, user friendliness will become even more important, and we expect the situation-oriented approach to be of even greater value with the new robot than before.

References

Babcock, P. (1976): Webster’s Third New International Dictionary of the English Language, G. & C. Merriam Company, Springfield, MA, U.S.A.

Brooks, R. A. (1996): From Earwigs to Humans. To appear: Journal of Robotics and Autonomous Systems.

Graefe, V. (1992): Vision for Autonomous Mobile Robots. Proceedings, IEEE Workshop on Advanced Motion Control. Nagoya, pp 57-64.

Graefe, V.; Wershofen, K. P. (1991): Robot Navigation and Environmental Modeling. International Advanced Robotics Programme – Proceedings of the Second Workshop on Multi-sensor Fusion and Environment Modeling. Oxford.

Wershofen, K. P. (1996): Zur Navigation sehender mobiler Roboter in Wegenetzen von Gebäuden – Ein objektorientierter verhaltensbasierter Ansatz. Dissertation, Fakultät für Luft- und Raumfahrttechnik der Universität der Bundeswehr München.