

Situated Robotics

Maja J Matarić, University of Southern California, Los Angeles, CA, USA

CONTENTS

Introduction
Types of robot control

Comparison and discussion

Situated robotics is the study of robots embedded in complex, often dynamically changing environments. The complexity of the robot control problem is directly related to how unpredictable and unstable the environment is, to how quickly the robot must react to it, and to how complex the task is.

INTRODUCTION

Robotics, like any concept that has grown and evolved over time, has eluded a single, unifying definition. What once used to be thought of as a replacement for repetitive, manual labor, has grown into a large field that includes applications as diverse as automated car assembly, space exploration and robotic soccer. Although robotics includes teleoperation, in which the robot itself may be merely a remotely-operated body, in most interesting cases the system exists in the physical world, typically in ways involving movement. Situated robotics, focuses specifically on robots that are embedded in complex, challenging, often dynamically changing environments. 'Situatedness' refers to existing in, and having one's behavior strongly affected by such an environment. Examples of situated robots include autonomous robotic cars on the highway or on city streets (Pomerleau 1989), teams of interacting mobile robots (Mataric' 1995), a mobile robot in a museum full of people (Burgard et al, 2000). Examples of unsituated robots, which exist in fixed, unchanging environments, include assembly robots operating in highly structured, strongly predictable environments. The predictability

and stability of the environment largely determines the complexity of the robot that must exist in it; situated robots present a significant challenge for the designer.

Embodiment is a concept related to situatedness. It refers to having a physical body interacting with the environment through that body. Thus, embodiment is a form of situatedness: an agent operating within a body is situated within it, since the agent's actions are directly and strongly affected by it. Robots are embodied: they must possess a physical body in order to sense their environment, and act and move in it. Thus, in principle every robot is situated. But if the robot's body must exist in a complex, changing environment, the situatedness, and thus the control problem, are correspondingly complex.

TYPES OF ROBOT CONTROL

Robot control is the process of taking information about the environment, through the robot's sensors, processing it as necessary in order to make decisions about how to act, and then executing those actions in the environment. The complexity of the environment, i.e., the level of situatedness, clearly has a direct relation to the complexity of the control (which is directly related to the task of the robot): if the task requires the robot to react quickly yet intelligently in a dynamic, challenging environment, the control problem is very hard. If the robot need not respond quickly, the required complexity of control is

reduced. The amount of time the robot has to respond, which is directly related to its level of situatedness and its task, influences what kind of controller the robot will need.

While there are infinitely many possible robot control programs, there is a finite and small set of fundamentally different classes of robot control methodologies, usually embodied in specific robot control architectures. The four fundamental classes are: reactive control ('don't think, react'), deliberative control ('think, then act'), hybrid control ('think and act independently in parallel'), and behavior-based control ('think the way you act').

Each of the approaches above has its strengths and weaknesses, and all play important and successful roles in certain problems and applications. Different approaches are suitable for different levels of situatedness, the nature of the task, and the capabilities of the robot, in terms of both hardware and computation. Robot control involves the following unavoidable trade-offs:

Thinking is slow, but reaction must often be fast.

Thinking allows looking ahead (planning) to avoid bad actions. But thinking too long can be dangerous (e.g., falling off a cliff, being run over).

To think, the robot needs potentially a great deal of accurate information. Information must therefore actively be kept up to date. But the world keeps changing as the robot is thinking, so the longer it thinks, the more inaccurate its solutions.

Some robots do not 'think' at all, but just execute preprogrammed reactions, while others think a lot and act very little. Most lie between these two extremes, and many use both thinking and reaction. Let us review each of the four major approaches to robot control, in turn.

Reactive Control

'Don't think, react!' Reactive control is a technique for tightly coupling sensory inputs and effector outputs, to allow the robot to respond very quickly to changing and unstructured environments (Brooks, 1986). Reactive control is often described as its biological equivalent: 'stimulus-response'. This is a powerful control method: many animals are largely reactive. Thus, this is a popular approach to situated robot control. Its limitations include the robot's inability to keep much information, form internal representations of the world (Brooks 1991a), or learn over time. The tradeoff is made in favor of fast reaction time and against complexity of reasoning. Formal analysis has shown that for environments and tasks that can be characterized *a priori*, reactive controllers can be shown to be highly powerful, and, if properly structured, capable of optimal performance in particular classes of problems (Schoppers 1987; Agre and Chapman 1990). But in other types of environments and tasks, where internal models, memory, and learning are required, reactive control is not sufficient.

Deliberative Control

'Think, then act.' In deliberative control, the robot uses all of the available sensory information, and all of the internally stored knowledge, to reason about what actions to take next. The reasoning is typically in the form of planning, requiring a search of possible state-action sequences and their outcomes. Planning, a major component of artificial intelligence, is known to be a computationally complex problem. The robot must construct and then evaluate potentially all possible plans until it finds one that will tell it how to reach the goal, solve the problem, or decide on a trajectory to execute. Planning requires the existence of an internal representation of the world, which allows the robot to look ahead into the future, to predict, the outcomes of possible actions in various states, so as to generate plans. The

internal model, thus, must be kept accurate and up-to-date. When there is sufficient time to generate a plan, and the world model is accurate, this approach allows the robot to act strategically, selecting the best course of action for a given situation. However, being situated in a noisy, dynamic world usually makes this impossible. Thus, few situated robots are purely deliberative.

Hybrid Control

‘Think and act independently in parallel.’ Hybrid control combines the best aspects of reactive and deliberative control: it attempts to combine the real-time response of reactivity with the rationality and efficiency of deliberation. The control system contains both a reactive and a deliberative component, and these must interact in order to produce a coherent output. This is difficult: the reactive component deals with the robot's immediate needs, such as avoiding obstacles, and thus operates on a very short time-scale and uses direct external sensory data and signals; while the deliberative component uses highly abstracted, symbolic, internal representations of the world, and operates on a longer time-scale. As long as the outputs of the two components are not in conflict, the system requires no further coordination. However, the two parts of the system must interact if they are to benefit from each other. Thus, the reactive system must override the deliberative one if the world presents some unexpected and immediate challenge; and the deliberative component must inform the reactive one in order to guide the robot toward more efficient trajectories and goals. The interaction of the two parts of the system requires an intermediate component, whose construction is typically the greatest challenge of hybrid design. Thus, hybrid systems are often called ‘three layer systems’, consisting of the reactive, intermediate, and deliberative layers. A great deal of research has been conducted on how to designing these components and their interactions (Giralt *et*

al.,1983; Firby, 1987; Arkin, 1989; Malcolm and Smithers, T., 1990; Connell, 1991; Gat, 1992).

Behavior-Based Control

‘Think the way you act.’ Behavior-based control draws inspiration from biology, and tries to model how animals deal with their complex environments. The components of behavior-based systems are called behaviors: these are observable patterns of activity emerging from interactions between the robot and its environment. Such systems are constructed in a bottom-up fashion, starting with a set of survival behaviors, such as collision-avoidance, which couple sensory inputs to robot actions. Behaviors are added to provide more complex capabilities, such as wall following, target chasing, exploration, and homing. New behaviors are introduced into the system incrementally, from the simple to the more complex, until their interaction results in the desired overall capabilities of the robot. Like hybrid systems, behavior-based systems may be organized in layers, but unlike hybrid systems, the layers do not differ from each other greatly in terms of time-scale and representation used. All the layers are encoded as behaviors, processes that take inputs and send outputs to each other.

Behavior-based systems and reactive systems share some similar properties: both are built incrementally, from the bottom up, and consist of distributed modules. However, behavior-based systems are fundamentally more powerful, because they can store representations (Mataric, 1992), while reactive systems cannot do so. Representations in behavior-based systems are stored in a distributed fashion, so as to best match the underlying behavior structure that causes the robot to act. Thus if a robot needs to plan ahead, it does so in a network of communicating behaviors, rather than a single centralized planner. If a robot needs to store a large map, the map is likely to be distributed over multiple behavior modules representing its components, like a network of landmarks, as in

(Mataric, 1990), so that reasoning about the map can be done in an active fashion, for example using message passing within the landmark network. Thus, the planning and reasoning components of the behavior-based system use the same mechanisms as the sensing and action-oriented behaviors, and so operate on a similar time-scale and representation. In this sense, ‘thinking’ is organized in much the same way as ‘acting’.

Because of their capability to embed representation and plan, behavior-based control systems are not an instance of ‘behaviorism’ as the term is used in psychology: behaviorist models of animal cognition involved no internal representations. Some argue that behavior-based systems are more difficult to design than hybrid systems, because the designer must directly take advantage of the dynamics of interaction rather than minimize interactions through traditional system modularity. However, as the field is maturing, expertise in complex system design is growing, and principled methods of distributed modularity are becoming available, along with behavior libraries. Much research has been conducted in behavior-based robot control.

COMPARISON AND DISCUSSION

Behavior-based systems and hybrid systems have the same expressive and computational capabilities: both can store representations and look ahead. But they work in very different ways, and the two approaches have found different niches in mobile robotics problem and application domains. For example, hybrid systems dominate the domain of single robot control, unless the domain is so time-demanding that a reactive system must be used. Behavior-based systems dominate the domain of multi-robot control, because the notion of collections of behaviors within the system scales well to collections of such robots, resulting in robust, adaptive group behavior.

In many ways, the amount of time the robot has (or does not have) determines what type of controller will be most appropriate. Reactive systems are the best choice for environments

demanding very fast responses; this capability comes at the price of not looking into the past or the future. Reactive systems are also a popular choice in highly stochastic environments, and environments that can be properly characterized so as to be encoded in a reactive input-output mapping. Deliberative systems, on the other hand, are the best choice for domains that require a great deal of strategy and optimization, and in turn search and planning. Such domains, however, are not typical of situated robotics, but more so of scheduling, game playing, and system configuration, for instance. Hybrid systems are well suited for environments and tasks where internal models and planning can be employed, and the real-time demands are few, or sufficiently independent of the higher-level reasoning. Thus, these systems ‘think while they act.’ Behavior-based systems, in contrast, are best suited for environments with significant dynamic changes, where fast response and adaptivity are necessary, but the ability to do some looking ahead and avoid past mistakes is required. Those capabilities are spread over the active behaviors, using active representations if necessary (Mataric, 1997). Thus, these systems ‘think the way they act.’

We have largely treated the notion of ‘situated robotics’ here as a problem: the need for a robot to deal with a dynamic and challenging environment it is situated in. However, it has also come to mean a particular class of approaches to robot control, driven by the requirements of situatedness. These approaches are typically behavior-based, involving biologically-inspired, distributed, and scalable controllers that take advantage of a dynamic interaction with the environment rather than of explicit reasoning and planning. This overall body of work has included research and contributions in single-robot control for navigation (Connell, 1990; Mataric, 1990), models of biological systems ranging from sensors to drives to complete behavior patterns (Beer, 1990; Cliff, 1990; Maes, 1990; Webb, 1994; Blumberg, 1996), robot soccer (Asada *et al.*, 1994; Werger, 1999; Asada *et al.*, 1998),

cooperative robotics (Matarić, 1995; Kube, 1992; Krieger *et al.*, 2000; Gerkey and Matarić, 2000), and humanoid robotics (Brooks and Stein, 1994; Scassellati, 2000; Matarić, 2000). In all of these examples, the demands of being situated within a challenging environment while attempting to safely perform a task (ranging from survival, to achieving the goal, to winning a soccer match) present a set of challenges that require the robot controller to be real-time, adaptive, and robust.

The ability to improve performance over time, in the context of a changing and dynamic environment, is also an important area of research in situated robotics. Unlike in classical learning, where the goal is to optimize performance over a typically long period of time, in situated learning the aim is to adapt relatively quickly, achieving greater efficiency in the light of uncertainty. Models from biology are often considered, and reinforcement learning models are particularly popular, given their ability to learn directly from environmental feedback.

This area continues to expand and address increasingly complex robot control problems. There are several good surveys on situated robotics which provide more detail and references (e.g. Brooks, 1991b; Matarić, 1998).

References

- Agre P and Chapman D (1990) What are plans for? In: Maes P (ed) *Designing Autonomous Agents*, pp.17-34. Cambridge, MA: MIT Press.
- Arkin R (1989) Towards the unification of navigational planning and reactive control. In: *Proceedings, American Association for Artificial Intelligence Spring Symposium on Robot Navigation*, pp.1-5. Palo Alto, CA: AAAI/MIT Press.
- Asada M, Stone P, Kitano H *et al.* (1998) The RoboCup physical agent challenge: Phase I. *Applied Artificial Intelligence* **12**: 251-263.
- Asada M, Uchibe E, Noda S, Tawaratsumida S and Hosoda K (1994) Coordination of multiple behaviors acquired by a vision-based reinforcement learning. In: *Proceedings, IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, pp.917-924. Munich: IEEE Computer Society Press.
- Beer R, Chiel H and Sterling L (1990) A biological perspective on autonomous agent design. *Robotics and Autonomous Systems* **6**: 169-186.
- Blumberg B (1996) *Old Tricks, New Dogs: Ethology and Interactive Creatures*. PhD thesis, MIT.
- Brooks A (1991a) Intelligence without representation. *Artificial Intelligence* **47**: 139-160.
- Brooks A (1991b) Intelligence without reason. In: *Proceedings, International Joint Conference on Artificial Intelligence Sydney, Australia*, pp.569-595. Cambridge, MA. MIT Press.
- Brooks R (1986) A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* **2**: 14-23.
- Brooks R and Stein L (1994) Building brains for bodies. *Autonomous Robots* **1**: 7-25.
- Burgard W, Cremers A, Fox D *et al.* (2000) Experiences with an interactive museum tour-guide robot. *Artificial Intelligence* **114**: 32-149.
- Cliff D (1990) The computational hoverfly; a study in computational neuroethology. In: Meyer J-A and Wilson S (eds) *Proceedings, Simulation of Adaptive Behavior*, pp. 87-96. Cambridge, MA: MIT Press.
- Connell J (1990) *Minimalist Mobile Robotics: A Colony Architecture for an Artificial Creature*. Boston, MA: Academic Press.
- Connell J (1991) SSS: a hybrid architecture applied to robot navigation. In: *Proceedings, International Conference on Robotics and Automation, Nice, France*, pp. 2719-2724. Los Alamitos, CA: AAAI/MIT Press.
- Firby J (1987) An investigation into reactive planning in complex domains. In: *Proceedings of the Sixth National Conference of the American Association for Artificial Intelligence Conference*, pp. 202-206 Seattle, WA: AAAI/MIT Press.
- Gat E (1998) On three-layer architectures. In: Kortenkamp D, Bonnasso R and Murphy R (eds) *Artificial Intelligence and Mobile Robotics*. AAAI Press.
- Gerkey B and Matarić M (2002) Principled communication for dynamic multi-robot task allocation. In: Rus D and Singh S (eds)

- Proceedings of the International Symposium on Experimental Robotics 2000, Waikiki, Hawaii*, pp. 341-352. Berlin Heidelberg: Springer-Verlag.
- Giralt G, Chatila R and Vaisset M (1983) An integrated navigation and motion control system for autonomous multisensory mobile robots. In: *Proceedings of the First International Symposium on Robotics Research*, pp. 191-214. Cambridge, MA: MIT Press.
- Krieger M, Billeter J-B and Keller L (2000) Ant-like task allocation and recruitment in cooperative robots. *Nature* **406**: 992.
- Kube R and Zhang H (1992) Collective robotic intelligence. In: *Proceedings, Simulation of Adaptive Behavior*, pp. 460-468. Cambridge, MA: MIT Press.
- Maes P (1990) Situated agents can have goals. *Robotics and Autonomous Systems* **6**: 49-70.
- Malcolm C and Smithers T (1990) Symbol grounding via a hybrid architecture in an autonomous assembly system. *Robotics and Autonomous Systems* **6**: 145-168.
- Matarić M (1990) Navigating with a rat brain: a neurobiologically-inspired model for robot spatial representation. In: Meyer J-A and Wilson S (eds) *Proceedings, From Animals to Animats 1, First International Conference on Simulation of Adaptive Behavior*. pp. 169-175. Cambridge, MA: MIT Press.
- Matarić M (1992) Integration of representation into goal-driven behavior-based robots. *IEEE Transactions on Robotics and Automation* **8** (3): 304-312.
- Matarić M (1995) Designing and understanding adaptive group behavior. *Adaptive Behavior* **4**(1): 51-80.
- Matarić M (1997) Behavior-based control: examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence* **9**: 323-336.
- Matarić M (1998) Behavior-based robotics as a tool for synthesis of artificial behavior and analysis of natural behavior. *Trends in Cognitive Science* **2**(3): 82-87.
- Matarić M (2000) Getting humanoids to move and imitate. *IEEE Intelligent Systems* **15**(4): 18-24.
- Pomerleau D (1989) ALVINN: an autonomous land vehicle in a neural network. In: Touretzky D (ed) *Advances in Neural Information Processing Systems 1*, pp. 305-313. San Mateo, CA: Morgan Kaufmann.
- Scassellati B (2001) Investigating models of social development using a humanoid robot. In Webb B and Consi T (eds) *Biorobotics*, pp.145-168. Cambridge, MA: MIT Press.
- Schoppers M (1987) Universal plans for reactive robots in unpredictable domains. In: *Proceedings, IJCAI-87*, pp. 1039-1046. Menlo Park, CA: Morgan Kaufman
- Webb B (1994) Robotic experiments in cricket phonotaxis. In: *Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, pp. 45-54. Cambridge, CA: MIT Press.
- Werger B (1999) Cooperation without deliberation: a minimal behavior-based approach to multi-robot teams. *Artificial Intelligence* **110**: 293-320.

Further Readings

- Arkin R (1998) *Behavior-Based Robotics*. Cambridge, MA: MIT Press.
- Brooks R (1999) *Cambrian Intelligence*. Cambridge, MA: MIT Press.
- Maes P (1994) Modeling adaptive autonomous agents. *Artificial Life* **2**(2): 135-162.
- Russell S and Norvig P (1995) *Artificial Intelligence. A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall.

Glossary

- Autonomous robot** A robot capable of performing without any external user or operator intervention.
- Behavior-based robot control** Using collections of behaviors (which may be reactive or may contain state and internal representations) to structure robot control.
- Deliberative robot control** The use of centralized representations and planning methods for generating a sequence of actions for the robot to perform.

Embodiment A form of situatedness, having a body and having one's actions directly and strongly affected and constrained by that body.

Hybrid robot control Using a combination of methods, typically a combination of deliberative and reactive control, to control a robot.

Learning robots Robots capable of improving their performance over time, based on past experience.

Reactive robot control The use of only reactive rules, and no internal memory or planning, in order to enable the robot to quickly react to its environment and task.

Robot A physical system equipped with sensors (e.g., cameras, whiskers, microphones, sonars) and effectors (e.g., arms, legs, wheels) that

takes sensory inputs from its environment, processes them, and acts on its environment through its effectors in order to achieve a set of goals.

Robot control The process of taking information about the environment, through the robot's sensors, processing it as necessary in order to make decisions about how to act, and then executing those actions in the environment.

Situated robotics The field of research that focuses on robots that are embedded in complex, challenging, often dynamically changing environments.

Situatedness Existing in, and having one's behavior strongly affected by a complex environment.

Keywords: (Check)

Robotics; situatedness; embodiment; learning; autonomy