

Integration of Machine Learning and Sensor-Based Control in Intelligent Robotic Systems *

Gerald DeJong Seth Hutchinson Mark W. Spong
University of Illinois
Urbana, Illinois

Abstract

This paper discusses the integration of machine learning and sensor-based control in intelligent robotic systems. Our research is interdisciplinary and combines techniques of explanation-based control with robust and adaptive nonlinear control, computer vision, and motion planning. Our intent in this research is to go beyond the strict hierarchical control architectures typically used in robotic systems by integrating modeling, dynamics, and control across traditional levels of planning and control at all levels of intelligence. Our ultimate goal is to combine analytical techniques of nonlinear dynamics and control with artificial intelligence into a single new paradigm in which symbolic reasoning holds an equal place with differential equation based modeling and control.

1 Introduction

The development of robots which are capable of complex, autonomous behavior such as adapting to changes and uncertainties in their environments, planning and executing strategies to carry out tasks without human intervention, and learning from past experience to improve future performance is one of the ultimate goals of robotics research. Achieving this goal would have tremendous impact in many areas of engineering. It is clear that many of the research problems that need to be solved to achieve this goal lie in the area broadly classified as intelligent control, and consist of modeling and control of uncertain nonlinear dynamical systems, and the integration of machine learning with sensor-based, real-time control.

Despite recent advances in robust and adaptive nonlinear control theory, present day robot manipulators and mobile robots are incapable of more than the most rudimentary of intelligent behaviors. Part of the reason for this is the fact that the intelligent execution and control of complex tasks by mechanical devices is difficult to accomplish within the traditional differential equation

based trajectory tracking paradigms of modern control theory. Trajectory generation and tracking control techniques are suitable for a limited class of tasks, such as spray painting, arc welding, or palletizing, but they are ill-suited to more complex tasks for which mathematical models are unavailable, poorly known, or too complex to be practical. Examples of complex tasks for which simple mathematical models and simple trajectories are not easily described include building a house, repairing a satellite, navigating through a network of pipes to inspect for corrosion, repairing downed power lines, cleaning up hazardous waste, and many others.

An *intelligent* robotic system would possess a symbolic reasoning capability and would thus be able to reason about the *qualitative* behavior of the system. While symbolic reasoning of this sort has been a mainstay of the Artificial Intelligence community for many years, control theorists have strongly criticized traditional symbolic AI paradigms of reasoning about action and change for their neglect of real-world effects faced by a robotic system, such as dynamic uncertainties, disturbances, noisy measurements, bandwidth limitations, etc. The classic examples of inserting a peg into a hole, and stacking blocks are applications where success or failure depends not only on reasoning about the sequence of motions necessary to accomplish the task but also upon understanding and controlling the interaction forces, the dynamics of the robot and its environment, etc. In spite of the obvious need and advantages of integrating planning and control, current robotic systems tend to be hierarchical instead, with high level task planners at the top which compute trajectories without utilizing information about the dynamics of the process, and low level servo systems at the bottom, which are designed simply to track the trajectories presented to them.

To date AI and vision research has failed to exploit the analytical results and techniques from control theory. The new paradigm that we will explore is based fundamentally on the belief that Intelligent Control of Robotic Systems is not possible without integrating the three areas of machine learning, sensing, and control theory. Some specific elements of this new paradigm

*This research is partially supported by The National Science Foundation and the Electric Power Research Institute under the joint NSF/EPRI Intelligent Control Initiative, Grant number ECS-9216428

are:

1. Planning must take into account what information the sensing system can robustly and quickly deliver. In particular, since general purpose vision currently is not possible in real-time, an intelligent control system must move beyond the traditional role of merely "observing" sensor data and take an active role in planning the sensing strategies to be used.
2. The sensing system should exploit techniques from control theory to aid in the acquisition and processing of sensor data. For example, an important problem in current robotics research is using a moveable camera to obtain vision feedback for end-effector servo control. This tracking problem could be greatly aided by the use of control theoretic methods, both for image plane feature tracking, and for controlling the motions of the tracking camera. Sensing is thus a key element to bridge the AI—Control Theory gap.
3. Planning must take into account the performance limits of the control system. For example, one of the well-known ways that an adaptively controlled system can become unstable is if the reference trajectory excites the unmodeled dynamics. Thus a hierarchical system in which planning takes place independently of control can fail dramatically if adaptive control is used in an attempt to increase its performance.
4. Control techniques must be developed to cope with uncertainties in the system beyond the usual cases of unknown parameters and unknown high frequency components. To this end, symbolic techniques can be used to diagnose model inadequacies and to suggest appropriate refinements in response to observed responses.
5. Machine learning techniques must be developed to suggest general descriptions of systematic behavior underlying world observations. This capability is essential if intelligent robotic systems are to perform effectively in novel situations, or in environments that change dynamically in unexpected ways.

The intent of our research is to increase the scope of control theory to embrace applications which cannot be solved with classical, modern, fuzzy logic, or neural net control systems. In particular, we believe that the AI area of planning (reasoning about perception, action, and world change) and the discipline of control theory (modeling, sensing, and actuation) are designed to serve the same purpose. In both, the task is to construct a strategy to transform a system from a given (perhaps unknown) initial state to a desired final state.

2 Task Control

Task Control involves the integration of control theory and AI planning into a new synthesis. We focus on hybrid systems for which no full adequate mathematical model can be given. An illustrative example is a

robot working in a complex environment in which friction, inelastic collisions, and other uncertainties dominate. Traditionally, robot control is viewed as separate from trajectory generation, grasping, etc. which are relegated to a separate planner. A more natural viewpoint, however, is to consider the robot, for which an adequate model can be constructed, and the robot's environment, which cannot be adequately described mathematically, as a single integrated system. In fact, such a viewpoint is taken by most of control theory (excluding robotics). One does not, for example, imagine the position of a valve in a chemical plant to be the output of the control system. One typically wishes to control the level of a chemical, not the position of a valve whose effect varies with other factors. No control engineer would say "Ah, well, that's a planning problem. I only control actuators." Yet that is just the response given by robotics control engineers.

Why is robotics different? Why not view the robot-environment as a single plant? One reason might be that it seems to make a convenient dividing line. This is traceable to the "general-purposeness" of robots. The chemical plant actuator will never find itself in a different environment. It always controls a particular flow of a particular fluid. Plant engineers are not allowed to redo the plumbing on a daily or hourly schedule so that at one moment the valve controls the flow of a high pressure liquid of a certain density and viscosity and another time a low pressure gas. In non-robotic control, the changes in environment are simple and can be circumscribed and modeled. This cannot be said of robots, which by their very nature are general purpose. A robot may be at one time quickly moving a heavy object grasped far from its center of mass in which little path accuracy is required, at another time applying a force to an unmoving tower to steady it, and still later precisely positioning a small part. The environment changes so drastically and so quickly that attempting to analyze "the environment" as a single plant is ludicrous from a control theory point of view.

3 Machine Learning

Recently, a new technique called explanation-based control [3] has been advanced. It appears to be a promising vehicle to bridge the current gap between the symbolic reasoning methods that underlie AI planning and the existence of continuous changes in the world, a cornerstone of control theory.

Explanation-based control employs machine learning techniques, primarily explanation-based learning [4,10] over a symbolic axiom set representing background knowledge of the world. The ontology of the symbolic logic is inspired by control theory and permits the representation of simultaneously and continuously varying world quantities. The predicates are based on work in qualitative reasoning [5,2,8,14]. Basically, the approach involves 1) observing a human or specialized control system expert as it solves a problem currently beyond the

learning system's capabilities, 2) constructing a symbolic qualitative explanation for why the expert's behavior results in the desired profile of effects, 3) symbolically generalizing the explanation in standard EGGs fashion [11], 4) calibrating the concept with the quantitative points observed from the expert's behavior, 5) using numerical interpolation between observed points to efficiently estimate the world's behavior, 6) quantitatively refining the planning concept with additional observed world points derived from the concept's use, 7) symbolically refining the concept by conjecturing an alternative explanation when world observations contradict the current explanation.

Through symbolic explanation (which identifies relevant inputs and state variables) coupled with numeric interpolation (to generalize the expected world behavior beyond the observed sample points) an explicit empirically derived numeric approximation of the world's inverse dynamics is formed. These explicit functions can be employed in control applications. The crucial point is that no differential equation model of the world's behavior is used. Rather, the system relies upon a symbolic axiom set describing the world's qualitative behavior. In many applications for which an adequate quantitative differential equation model of the world is too complex or unknown, one may still be able to provide an adequate qualitative description in symbolic terms. Furthermore, preliminary experience with the approach indicates that it may be tolerant of some non-linearities in the world, and may exhibit a certain robustness without the usual concomitant sluggishness.

Some form of stability analysis is most crucial; an analog to controllability and observability are also desirable. These are the primary research issues we intend to pursue for explanation-based control.

4 Adaptive Control

In the past decade important strides have been made in the design and analysis of robust and adaptive controllers for robotic manipulators using various mathematical tools, such as Lagrangian and Hamiltonian mechanics, passive network synthesis, Lyapunov methods, differential geometric control theory, singular perturbations and integral manifolds. The design and analysis of Intelligent Controllers requires the broadening of such tools to encompass real-time vision feedback, machine learning and AI planning techniques. This is a nontrivial task. To begin with, adaptive control in conjunction with real-time vision feedback is a completely open area of research in robotics. Adaptive control is attractive because it holds out the promise of improved performance over a wide range of payloads. However, until the robustness properties of adaptive robot control are fully understood, industrial designers will be reluctant, even unwise, to use them. Given the complicated behavior that can arise from adaptive control of even a simple first order linear system[9], it should not be surprising that a complete understanding of the robustness

of adaptive control of nonlinear systems as complicated as multi-link, multi-sensor robotic manipulators is still lacking. With the inclusion of vision feedback in the control loop the problem of robustness of adaptive control becomes even more complicated. With the further integration of explanation based learning directly into the feedback loops, perhaps operating in multiple time scales, we are faced with a class of nonlinear systems for which few prior results are available to help in the analysis or design.

It is known that the stability of adaptive systems are highly sensitive to disturbances and unmodeled dynamics. These arise in the robotics context from several sources. External disturbances include many types of interaction with the environment. For example, robotic assembly has been described as a sequence of controlled collisions with the environment. These collision forces can be viewed as disturbances to the controller. A repetitive task, for example, subjects the robot to periodic forcing which, even in non-adaptive control, can excite complex nonlinear dynamic behavior, such as period doubling bifurcations and chaos[15].

Unmodeled dynamics include actuator/sensor dynamics, joint flexibility, link flexibility, and environment dynamics. Environment dynamics arise in force and impedance controlled tasks such as assembly and grinding and will become increasingly important in future applications.

Several so-called "instability mechanisms" in adaptive control have been identified [7]. Among the mechanisms leading to instability are:

- 1) Reference trajectories which are "too fast." If the bandwidth of the reference trajectory is in the same frequency range as the unmodeled dynamics, then these unmodeled dynamic modes can be excited, leading to instability.
- 2) Parameter drift. In typical parameter adaptive controllers the estimated parameters are not guaranteed to converge to their true values, but only to remain bounded, without additional persistency of excitation conditions.
- 3) High gain instability. This type of instability, when the controller gains excite unmodeled dynamics, is actually due to a loss of passivity from the rigid robot case and can occur even for nonadaptive algorithms [16].
- 4) Fast adaptation instability. This type of instability occurs when the gains in the parameter update law are too large.
- 5) Neglected time delays. It has been known since the mid 1960's that small time delays in data transmission in systems such as bilateral force reflecting teleoperators can be destabilizing.

In addition to these, there are a number of instability mechanisms that are directly attributable to the introduction of vision sensing into the control loop.

- 6) **Lens Distortion.** The geometric correspondence between points in the 3D space and points in an image is typically approximated by either perspective projection (which reflects an ideal thin lens assumption) or by orthographic projection (which holds only when the distance from the camera to the object is much greater than the lens focal length). Neither of these approximations account for limited depth of field, vignetting, and non-linear distortions of the image as the radial distance from the focal center increases.
- 7) **Image Plane Quantization Effects.** The camera image plane is actually a discrete array of sensors. Therefore, the coordinates of object features used for visual tracking are actually quantized versions of the true coordinates of those features.
- 8) **Motion Induced Blur in Image Formation.** If there is relative motion between the camera and the object being imaged, then the resulting image will be blurred.
- 9) **Low Frequency Vision Sampling.** To date, the fastest sampling rate that has been reported in the literature on visual servo control is still well below video rate. Such slow sampling rates have severe implications for system stability.

5 Experimental Research

We are developing a testbed system that can be used as an experimental platform to test our theoretical results. The testbed will be built around a planar, three-degree-of-freedom robot arm equipped with a wrist force/torque sensor and a real-time vision system. The camera may either be mounted away from the arm or eye-in-hand, depending on the research problem being investigated.

We will focus initially on *Robotic Air Hockey*, i.e., controlling a puck sliding on a surface, as the application to test our theoretical results. This application contains many of the fundamental research issues in intelligent control that are of interest to us. For example, the motion trajectory of the arm cannot be planned off-line and presented to the control system in a hierarchical fashion. Real-time planning in conjunction with real-time vision are an absolute necessity for this application. In addition, this is an application where an accurate model of the dynamics of the environment is difficult to obtain because of large uncertainty in friction combined with high sensitivity of the puck motion to initial conditions. To be more specific, one would expect large amounts of uncertainty and variation over time in the friction characteristics of the table and puck. In addition, irregularities in the table surface and bumpers would make it difficult to obtain a mathematical model of the environment sufficiently precise to enable accurate predictions of puck motion. When this large uncertainty is combined with the expected high sensitivity of the puck motion to changes in initial imposed forces and velocities, it appears that traditional motion planning and

tracking control schemes would be difficult or impossible to apply successfully. On the other hand, this application appears eminently suited to learning through repeated trials in combination with real-time adaptive vision feedback control.

Another area where this experimental set-up will facilitate cross-disciplinary research is in the nature and control of the interaction forces between the puck and the robot. Of course, it is only through the forces imparted to the puck during collision with the paddle held by the robot that the puck is controlled. Issues of force feedback control are thus of prime importance in this research. In this application, however, the time of contact between the puck and paddle are of extremely short duration, and force control, as it is traditionally conceived, has little meaning. Instead, the force information collected during impact must be combined with the information about the resulting puck motion as determined by the vision system and the information about the arm motion as determined by the joint encoders to modify the arm motion during repeated impacts. Standard learning control methods are not likely to succeed, however, because it is impossible to recreate the puck trajectory over two successive trials.

In our previous research we have developed and exploited the so-called network approach to force control[1] including the hybrid impedance control framework. In the hybrid impedance control framework we are able to control both force levels (i.e., track force trajectories) and simultaneously control the end-effector impedance. We may be able to apply these same techniques in the present context. For example, by repeated trials the robot may be able to learn the optimal impedance to use while imparting a desired force to the puck. Ultimately, such knowledge will pay tremendous benefits in numerous industrial settings, such as in robotic assembly applications, or robotic handling of radioactive waste or fuel rods in a nuclear power plant.

There are many interesting problems in robust and adaptive nonlinear control that arise in this example. The dynamics of this process are complex enough that they cannot simply be ignored in the learning and planning phase. For example, the intermittent, impulsive, and repetitive nature of the interaction forces, when combined with adaptive control may induce interesting and complex behaviors, such as bifurcations, chaotic motion, and bursting. Thus the control designer must be aware of the potential for such nonlinear phenomena and how the actions of motion planning and learning may excite them.

The *Robotic Air Hockey* application contains numerous interesting and fundamental problems in computer vision. For example, understanding the physics of puck motion in the image plane is necessary for accurate tracking and for commanding motion of the arm. Issues such as camera placement to maximize the visual analog of manipulability [16] need to be understood. In other words, it is desirable that "small puck motion" translate

to "large image motion" in order to increase the sensitivity of the vision sensor. However, it is also desirable that "large puck motion" translate to "small image motion" in order to reduce the amount of information that the vision system needs to process in real-time. Understanding this and other trade-offs are fundamental for this application.

References

- [1] Anderson, R.J., and Spong, M.W., "Hybrid Impedance Control of Robots, *IEEE J. of Robotics and Automation*, Vol. 4, no. 5, pp.549-556, Oct. 1988.
- [2] D. Bobrow (ed.), *Qualitative Reasoning about Physical System*. MIT Press, Boston, 1985.
- [3] G. F. DeJong, "Explanation-based control: An approach to reactive planning in continuous domains," In *Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling and Control*, San Diego, CA, November 1990.
- [4] G. F. DeJong and R. J. Mooney, "Explanation-based learning: An alternative view," *Machine Learning*, 1(2):145-176, 1986.
- [5] K. D. Forbus, "Qualitative process theory," *Artificial Intelligence*, 24:85-168, 1984.
- [6] Ghorbel, F., Fitzmorris, A., and Spong, M.W., "Robustness of Adaptive Control of Robot Manipulators: Theory and Experiment," *Workshop on Nonlinear and Adaptive Control: Application to Robotics*, Grenoble, France, Nov. 21-23, 1990, invited plenary paper.
- [7] Ioannou, P.A., and Kokotovic, P.V., "Instability Analysis and Improvement of Robustness of Adaptive Control," *Automatica*, 20 (1984), pp.583-594.
- [8] B. J. Kuipers, "Qualitative simulation," *Artificial Intelligence*, 29:289-338, 1986.
- [9] Middleton, R.H., and Kokotovic, P.V., "Boundedness Properties of Simple Indirect Adaptive Control Systems," *IEEE Trans. on Automatic Cont.*, Vol. AC-37, No. 12, pp. 1989-1994, Dec. 1992.
- [10] Thomas M. Mitchell, Richard M. Keller, and Smadar T. Kedar-Cabelli, "Explanation-based generalization: A unifying view," *Machine Learning*, 1(1):47-80, January 1986.
- [11] R. J. Mooney and S. W. Bennett, "A domain-independent explanation-based generalizer," In *Proceedings of the National Conference on Artificial Intelligence*, pages 551-555, Philadelphia, PA, August 1986.
- [12] Ortega, R., and Spong, M.W., "Adaptive Motion Control of Rigid Robots: A Tutorial," *Proc. IEEE Conference on Decision and Control*, Austin, TX, Dec. 1988.
- [13] Reed, J.S., and Ioannou, P.A., "Instability Analysis and Robust Adaptive Control of Robotic Manipulators," *Proc. 27th IEEE Conf. on Decision and Cont.*, Austin, Dec. 1988.
- [14] D. Weld and J. deKleer (ed.), *Readings in qualitative reasoning about physical systems*. Morgan Kaufmann, 1990.
- [15] Ydstie, B.E., and Golden, M.P., "Chaos and Strange Attractors in Adaptive Control Systems," *Proc. IFAC World Congress on Automatic Control*, Munich, July, 1987.
- [16] Yoshikawa, T., "Manipulability of Robotic Mechanisms," *The International Journal of Robotics Research*, Vol. 4, No. 2, pp. 3-9, Summer, 1985.