

# Building Robots

*Notes for a talk*

Alan K. Mackworth\*

Department of Computer Science  
University of British Columbia  
Vancouver, B. C., Canada V6T 1Z2

An intelligent agent prospers by maintaining implicit and explicit models of the environment and itself. Perception, reasoning and action must all serve the common purpose. Recent attempts, within artificial intelligence, to establish each of those areas as semi-autonomous disciplines have yielded useful mathematical and computational results but have also led to sterility. That strategy has failed to produce the coherent analytical science necessary for the synthetic engineering activity of building intelligent agents. Unlike Gaul, intelligence is not divisible into three parts. As Brooks points out, perception, reasoning and action do not correspond to natural scientific domains with clean interfaces and limited interaction amongst them; they correspond only to labels that we use to caricature aspects of the agent's behavior. But, although this reduction does not carry through, that's no excuse for abandoning reductionist scientific activity and retreating to holistic philosophizing. Alternate reductionist strategies are available, such as focussing on hierarchies of behavior units, each of which can embody elements of perception, reasoning and action, as in Brooks' "subsumption" architecture. Zhang and I have proposed the Constraint Nets model of intelligent systems as an alternative decomposition strategy that allows formal characterization and implementation techniques.

Neither AI nor robotics (nor, for that matter, computational vision or any other subdiscipline of either field) can proceed autonomously. Divide and conquer or, at least, the version of that game we have been playing, is not now the best strategy. The most payoff in the next few years will come from approaches that design, analyze and build integrated agents. This requirement for *cognitive integration*, the tight coupling of perception, reasoning and action, should dominate our research strategy.

Our idealizations and simple worlds lead us astray. In AI we postulated worlds in which all the effects of

an action are knowable before the action is taken in the world. We can characterize this as the Omniscient Fortune Teller Assumption (OFTA). The OFTA dictated, by *fiat*, that perception was unnecessary for intelligent action. It allowed an agent to retreat into its head, constructing a straight-line plan isomorphic to an action sequence which was then played as a motor command tape. The OFTA is now being relaxed (see, for example, the work on reactive planning) but it still permeates the way we design our agents. It has sanctioned the divorce of reasoning from perception and action. There is an interesting analogy here with motor control in robotics. The offline approach to planning is directly analogous to open loop dead reckoning control. They both embody the assumption of perfect knowledge of the consequences of all actions. The OFTA, and not the frame problem which follows from it, is the real difficulty here.

On the other hand, feedback control theory, using the perceived effects of actions to control future actions in order to achieve a desired purpose, has led to an array of mathematical and engineering triumphs. Moreover, hierarchical feedback control theory has shown us how to achieve stable behaviors for a wide variety of complex systems, by closing feedback loops between the agent and the world at every level of the hierarchical structure. This is achieved despite the stubborn reality of phenomena, such as joint backlash, friction and flexible links, that we cannot hope to model tractably. So far, however, hierarchical feedback control has mostly been used to control agents where the environmental description is impoverished: an  $n$ -dimensional vector of scalars. We need to apply the key insight of hierarchical feedback control but use descriptively richer languages and methodology from AI to model the environment and the agent itself.

By abandoning the OFTA, we see that the agent cannot maintain a faithful world model by reasoning alone. (From this it does not follow, *pace* Brooks, that we should abandon reason or representation!) Indeed, it cannot maintain a completely faithful world model by any means. Actions have many possible unpredictable out-

---

\* Shell Canada Fellow, Canadian Institute for Advanced Research. This work is supported, in part, by the Natural Sciences and Engineering Research Council of Canada and by the Institute for Robotics and Intelligent Systems Network of Centres of Excellence.

comes and real worlds cannot be exhaustively modelled. But, ranges and likelihoods of outcomes can be characterized and real worlds can be partially modelled. Uncertainty and risk-taking are a necessary component of intelligent behavior. Perception cannot be exhaustive; it must be purposive, model-based, incremental and multi-sensory. Perceptual actions can be planned and carried out to acquire knowledge. A blind person's cane tapping strategy illustrates the coupling of perception, reasoning and action: each subserves the others. Plans are programs. Straight-line code is only their simplest form. However, we must learn the automatic programming lesson. Even in the predictable, disembodied world inside a computer, automatic programming has proven an elusive goal. Automatic planning in the world of a robot is much harder.

AI and robotics will be integrated only if AI workers stop focussing on disembodied, solipsistic reasoners and if roboticists accept the need for richer, more adequate methodologies to describe the world. Non-standard (mostly non-deductive) logical approaches based on theory formation, dialectical reasoning, argument structures, belief as defeasible knowledge, situated automata and constraint-based model-theoretic approaches (as advocated by Poole, Genesereth, Shoham, Rosenschein, Reiter, Mackworth *et al.*) are all promising but they must consider perception and action as playing roles in the theory beyond simply providing truth values for atomic propositions. Overthrow the tyrannical reasoner! For example, Reiter and I have provided a logical framework for depiction that allows reasoning about a world and images of that world, characterizing the interpretations of an image as the logical models of the description of the image, the scene and the image-scene mapping. This allows the coupling of perception and reasoning through a common logic-based language.

The choice of target problem domain is key. It must require for its solution cognitive integration. It should require experimental and theoretical progress in techniques for perception, reasoning, and action but be within their grasp, so to speak. It should be useful with objective criteria for success, perhaps competing with another baseline technology. It should allow us to acknowledge the difficulty of automatic planning. Given all that, it should also be as simple, and exciting, as possible.

A target domain with these characteristics is telerobotics. Telerobotics is a further development beyond teleoperation. In teleoperation a human controls some remote device in a master-slave relationship. Telerobotics incorporates some autonomous robotic control with high-level human supervision. Such a system should have an

internal model of the environment and a model of itself. Mulligan, Lawrence and I have designed and built a model-based vision system that allows a telerobot to see and monitor its own limbs, allowing us to supplement or, perhaps, replace traditional joint sensors for position control. As the robot moves its limbs the perceptual system uses visual information and other senses to provide updates to its internal self-model. A typical hand-eye system has to hide its arm before looking at the scene. Surely one of the first perceptual tasks for a robot or a telerobot must be to understand images of its own moving body parts. Once it has achieved that, then visually-guided grasping and coordinated manipulation become possible. It also suggests using visual feedback to supplement or replace the traditional inverse kinematic and setpoint methods for path planning and path following which, again, can be seen as an extension of the offline planning method for robot action.

What we have done may be seen as a step towards achieving one of the goals set out earlier, namely, integrating control-theoretic and knowledge-based approaches. A robot manipulator is typically controlled by representing its configuration as a vector of joint angles. Individual servo loops for each joint allow precise control of the manipulator. In our model-based vision systems we are using an articulated, 3D model of the limb, a richer description than a vector of joint angles, to represent the proximal environment. But we envision using the perceptual data to close servo loops, allowing for the control of the movement of the limb continuously during an action. Another telerobotic environment we are experimenting with in our lab is based on elaborations of radio-controlled vehicles.

This approach achieves the necessary tight coupling of perception, reasoning and action. As specified above, the system is purposive, model-based, incremental and multisensory. Telerobotics, as an integrating application domain, has the advantage over building completely autonomous robots in that we can incrementally automate aspects of the total system's behavior while maintaining functionality. This gives us a common framework for the design of systems for a spectrum of applications ranging from human-controlled manipulators operating in constrained environments to autonomous agents in less structured environments. An agent's behavior must be specified and controlled at many levels: for example, at the joint level, at the end effector level and at the task level. At the lower levels that specification is in terms of set points and parameter vectors, at the higher levels as symbolic task descriptions. There are operational criteria for success: we cannot finesse reality by hiding in the OFTA. In order to satisfy those criteria, it must achieve cognitive integration.