

On Seeing Robots

Alan K. Mackworth*
Department of Computer Science
University of British Columbia
Vancouver, B.C.
Canada V6T 1W5

Abstract

Good Old Fashioned Artificial Intelligence and Robotics (GOFAIR) relies on a set of restrictive Omniscient Fortune Teller Assumptions about the agent, the world and their relationship. The emerging Situated Agent paradigm is challenging GOFAIR by grounding the agent in space and time, relaxing some of those assumptions, proposing new architectures and integrating perception, reasoning and action in behavioral modules. GOFAIR is typically forced to adopt a hybrid architecture for integrating signal-based and symbol-based approaches because of the inherent mismatch between the corresponding on-line and off-line computational models. It is argued that Situated Agents should be designed using a unitary on-line computational model. The Constraint Net model of Zhang and Mackworth satisfies that requirement. Two systems for situated perception built in our laboratory are described to illustrate the new approach: one for visual monitoring of a robot's arm, the other for real-time visual control of multiple robots competing and cooperating in a dynamic world.

1 Introduction

The title of this paper, “On Seeing Robots”, leaves substantial scope for playful exploration. The simple ambiguity is, of course, between describing robots that see their worlds and systems that see robots. These categories are not exclusive: I also combine them and discuss robots that see robots and even robots that see themselves. Furthermore, the title is designed to echo, and pay homage to, a classic vision paper entitled “On Seeing Things” by Max Clowes [1] as I have done once before [2]. But the context, the arguments and the conclusions are new; the comparison is used explicitly here to show the difference between the classical approach and an emerging situated approach to robotic perception. The most important reading of the title is that the paper is about how *we* see robots; it is about the computational paradigms, the assumptions, the architectures and the tools we use to design and build robots.

* Shell Canada Fellow, Canadian Institute for Advanced Research

2 Good Old Fashioned Artificial Intelligence and Robotics

The phrase Good Old Fashioned Artificial Intelligence (GOF AI) was introduced by Haugeland [3] to characterize the classical symbol manipulation approach to AI. In GOF AI intelligence is identified with reasoning and reasoning with rule-based manipulation of symbolic structures. Given the fact that syntactic proof theory and Tarskian semantic model theory can be placed in isomorphic correspondence, a GOF AI system can be said to reason about the real world. How it senses the world and how it acts in the world, if at all, are secondary concerns delegated to separate perception and action modules. We extend GOF AI here to Good Old Fashioned AI and Robotics (GOF AIR) to characterize the idea of building a robotic system with a perception front end that translates from signal to symbol, a GOF AI system as the meat in the sandwich and a motor back end that carries out actions in the world. So a GOF AIR system consists of three modules for perception, reasoning and action, respectively. (This characterization of a GOF AIR robot is, of course, an unfair but useful caricature.) The paradigmatic environment that a GOF AIR robot inhabits is the blocks world. Clowes [1] and many others [4] provided the tools to build perceptual systems that translated arbitrary images of that world to symbolic descriptions for the purposes of reasoning and planning. Planning for a GOF AIR robot, using the situation calculus or the simplified STRIPS representation, models actions as changes to a global world model, maintained as a set of sentences, to produce a plan. In GOF AIR (but not in general as we shall see) a plan is just a list of actions which if executed would change the world into its desired state, provided that the world were as modelled, the action models were correct and that nothing else intervened. It is possible to make explicit some of the meta-assumptions about the agent and its world implicit in much of the GOF AIR research strategy [5]:

- **Assumption IR (Individuals and Relations):** All that is useful for an agent can be described in terms of individuals and relations amongst individuals.
- **Assumption BK (Belief is Knowledge):** An agent's beliefs about the world are true and justified.
- **Assumption DK (Definite Knowledge):** An agent's knowledge of the world is definite and positive.
- **Assumption CK (Complete Knowledge):** The agent's knowledge of the world is complete. This requires that everything relevant about the world be known to the agent. This Closed World Assumption allows the agent to assume safely that a fact is false if it cannot infer that it is true.
- **Assumption SE (Static Environment):** The environment is static unless an agent changes it.
- **Assumption OA (One Agent):** There is only one agent in the world.

- **Assumption DW (Deterministic World):** Given a complete and definite description of the world the agent can predict all the effects of an action.
- **Assumption DSA (Discrete Sequential Actions):** Actions are discrete and they are carried out sequentially.

These assumptions are very restrictive. OA rules out other agents acting cooperatively to help the agent, competitively to frustrate the agent's plans or neutrally, as nature might do. OA also means that the agent does not have to react in real time to changes in the world. DW rules out non-deterministic actions, such as tossing a coin. BK, DK and CK mean that the agent is really omniscient — it has definite knowledge of everything relevant to achieving its goals. Assumption DSA rules out the need to consider continuous events such as processes, and the possibility of performing actions concurrently, which would require reasoning about the duration and termination of actions. By making all these assumptions explicit we can consider relaxing them independently, as needed.

To realize the force of these assumptions let us consider a world in which they are all violated. Suppose we want to build a robot to play soccer. Quite apart from all the difficult robotics and perception problems involved, we have substantial challenges in representation for planning and action. OA is violated: there are cooperating agents on the robot's team, competing agents on the other team, and neutral agents such as the referee and the weather. DW is violated: it is not possible to predict precisely where the ball will go when it is kicked, even if all the relevant factors are known. Each of BK, DK and CK is violated. Moreover, DSA is violated: continuous events such as a player running to a position, or the ball moving through the air, occur concurrently.

Our idealizations and simple worlds can lead us astray. The collective force of these assumptions is that, in GOFAIR, we postulate a world in which all the effects of an action are knowable before the action is taken in the world. In homage to this powerful consequence, we dub them the Omniscient Fortune Teller Assumptions (OFTA).

A further radical consequence of the OFTA is that they dictate that perception is unnecessary for intelligent action except as it is needed to determine the initial state of the world. They allow an agent to retreat into its head constructing, by reason alone, a plan as an action sequence which is then played as a motor command tape. In other words, planning is reduced to finding a straight-line program without conditionals or loops. Some of the OFTA are now being relaxed (see, for example, the work on reactive planning [6]) but they still permeate the way we design our agents. They have sanctioned the divorce of reasoning from perception and action. There is an interesting analogy here with motor control in robotics. The off-line approach to straight-line 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 them, is the real difficulty here. Just as dead reckoning fails for navigation, the unacceptable consequences of the OFTA have forced a crisis for GOFAIR which presages a paradigm shift. In the period of extraordinary science

provoked by the impending collapse of a paradigm there are many contenders for the new paradigm [7]. Some believe that a normal process of relaxing some of the OFTA assumptions will succeed; others that nothing short of a revolution will work. Either way, it is worth spending some time and effort to understand and make explicit the foundations of GOFAIR to see if they are all rotten or just a little shaky and in need of shoring up.

3 Situated Agents

The attempt, in the GOFAIR paradigm, to establish perception, reasoning and action as semi-autonomous disciplines has yielded useful mathematical and computational results but has 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. The perception, reasoning and action modules of GOFAIR not only can't be built but also do not correspond to natural scientific domains with clean interfaces and limited interaction amongst them. Perception, reasoning and action correspond only to labels that we use to caricature aspects of the agent's behavior. Brooks [8] has correctly pointed out that the traditional divide-and-conquer AI approach to robotics, by slicing intelligence into perception, reasoning and action, has pursued a strategy that does not scale up. This, incidentally, implies that any research program based on that division will be sterile. 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 the subsumption architecture [9]. I accept Brooks' diagnosis of the problem but not his prescription for the solution [10]. It is clear though that closer coupling of perception and action, intermediated by reasoning when necessary, in embedded behavioral modules is the correct general approach. As discussed later in this paper, an alternative decomposition strategy is the Constraint Net model of intelligent systems [11], 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. The version of divide-and-conquer that we have been playing, namely, functional decomposition, is not now the best strategy. The best 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. This is a non-trivial requirement: as I'll argue later, it follows as a consequence that systems must be designed and implemented in a single unitary framework.

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 [12] or representation [13]!) Indeed, it cannot maintain a completely faithful world model by any means. Actions have many possible unpredictable outcomes

and real worlds cannot be exhaustively modelled. But, ranges and likelihoods of outcomes can be characterized and real worlds can be partially modelled. Risk-taking under uncertainty is a necessary aspect of intelligent behavior. Perception is not exhaustive; it is purposive, model-based, situated, incremental and multi-modal. Perceptual actions are 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 robot 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. But planning, in its fully generality, is not a necessary component of an intelligent agent; however, responding appropriately to changes in the world is always necessary.

The claim is that AI and robotics will be integrated only if AI researchers stop focussing on disembodied, solipsistic reasoners and if roboticists accept the need for richer, more adequate methodologies to describe the world. Nonstandard logical approaches based on theory formation, dialectical reasoning, argument structures, belief as defeasible knowledge, situated automata and constraint-based model-theoretic approaches 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 Mackworth [14,15] 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 critiques and rejection, by some, of the GOFAIR paradigm have given rise to what we shall call the Situated Agent (SA) approaches of Rosenschein and Kaelbling, [16,17], Agre and Chapman [18,19], Smith [20], Brooks [12], Ballard [21], Winograd and Flores [22], Lavignon and Shoham [23], Zhang and Mackworth [24] and many others. The collection of SA approaches is sometimes also known loosely as Nouvelle AI. It is hard to define the SA approach succinctly; emerging paradigms can often only be defined in retrospect. Indeed, the various approaches hardly constitute a mutually consistent and coherent school; but, they do represent a movement. Perhaps a way to convey the flavor of the difference is that in GOFAIR *ad hoc* is a term of abuse (used, say, to describe a system without a Tarskian semantics); in SA, on the other hand, *ad hoc*, meaning literally "to this", is an indexical – a great compliment. In short, a situated agent is a real physical system grounded and embedded in a real world, here and now, acting and reacting in real-time.

Situated agents clearly indulge not only in situated action and, perhaps, in situated reasoning but also in situated perception [21,25]. Another shift in moving from GOFAIR to SA is from a single agent in a static world to multiple agents in a dynamic world

which, for our purposes, entails also a shift from static perception to dynamic perception. So one theme of this paper is *situated dynamic perception*.

Some of the connotations of the shift from GOFAIR to SA can be elicited by the shift from “Seeing Things” to “Seeing Robots”: the ultimately situated agent sees not randomly-arrayed, unexpected “things” but a coherent, dynamic evolving scene resulting, in part, from its own movements and actions. This shift is most dramatically and effectively conveyed when the robot sees parts of its own body.

4 Back to the Future

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 are hard 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 to model the environment and the agent itself.

Occam’s Razor requires that our most fundamental research goal should be to base the new paradigm on a unitary theory. Ideally such a theory will be mathematical in nature but will lead to appropriate computational formalisms. We already know that it must include standard control theory as a special case.

An alternative to a unitary theory is the approach, taken by many, of building hybrid systems with signal-based low-level systems and symbol-based high-level GOFAIR systems. The hybrid approach is esthetically repellent and pragmatically cumbersome; moreover, it has had limited experimental success.

The root problem with the hybrid approach is a complete mismatch of the nature of the two underlying computational paradigms [24]. The GOFAIR symbol-manipulating systems are based on *off-line* computational models such as virtual machines for Lisp or Prolog. In essence these are all in the off-line Turing Machine paradigm of computation. An off-line model computes its output as a mathematical *function* of its inputs. There is no notion that the inputs arrive over time. The signal-manipulating systems, though, are based on *on-line* models. An on-line model, such as a circuit, computes an output trace (a function of time, on a discrete, dense or event-based time structure, to a domain of values) as a *transduction* of its input traces. This fundamental mismatch ensures that the oft-discussed signal-symbol interface is hard, if not impossible, to specify coherently, let alone build.

Notice, in particular, that the off-line approach pervades GOFAIR. Planning, for example, is seen as an atemporal activity; it involves reasoning *about* actions in time but it does not occur *in* time. The recent flurry of activity in ‘anytime’ planning is an acknowledgment of this discrepancy. Vision is conceived as implementing a mathematical *function* whose input is the retinal stimulation and whose output is, variously, a description of the image, a viewer-centred description of the visible surfaces or a world-centred description. Deconstruction of GOFAIR along these lines is instructive, and perhaps necessary, if we are to escape the pervading off-line assumptions.

One of the requirements we place on a unitary paradigm is that it subsume, for example, signal processing, control systems, analog and digital circuit models, and dynamical systems, most generally. (This is indeed a tall order.) All of these paradigms assume an on-line computational model; they are also all of a venerable vintage. And yet the impression created by GOFAIR is that we have left these frameworks behind, or beneath, us. On the contrary, we must revisit them, include them and situate them in the symbolic paradigm; this requires substantial generalization of both the traditional signal-based and the traditional symbol-based approaches. (If this analysis is correct this move back to the future will indeed be ironic, and painful, both for GOFAIR and for Nouvelle AI; each is rather fond of thinking of itself as the *avant garde*.) The unitary approach will only succeed, following this line of argument, if that generalization is a single on-line computational model.

One such model is embodied in the Constraint Net (CN) approach that Ying Zhang and I have developed. CN is a model for robotic systems software implemented as modules with I/O ports [26]. A module performs a transduction from its input traces to its output traces, subject to the principle of causality: an output value at any time can depend only on the input values before, or at, that time. The language has a formal semantics based on the least fixpoint of sets of equations [11]. In applying it to a robot operating in a given environment one separately specifies the behaviour of the robot plant, the robot control program and the environment. The total system can then be shown to have various properties, such as safety and liveness, based on provable properties of its subsystems. This approach allows one to specify formally, and verify, models of embedded control systems. Our goal is to develop it as a practical tool for building real, complex, sensor-based robots. It can be seen as a development of Brooks’ subsumption architecture [8] that enhances its modular advantages while avoiding the limitations of the augmented finite state machine approach.

A robot situated in an environment is modeled as three machines: the robot plant, the robot control and the environment. Each is modeled separately as a dynamical system by specifying a CN with identified input and output ports. The robot is modeled as a CN consisting of a coupling of its plant CN and its control CN by identifying corresponding input and output ports. Similarly the robot CN is coupled to the environment CN to form a closed robot-environment CN.

The CN model is realized as an on-line distributed programming language with a formal algebraic denotational semantics and a specification language, a real-time temporal logic, that allows the designer to specify and prove properties of the situated robot by proving them of the robot-environment CN. So far, we have been able to specify, design, verify and implement systems for a robot that can track other robots [26], a robot that can escape from mazes and a two-handed robot that assembles objects [24], an elevator system [27] and a car-like robot that can plan and execute paths under non-holonomic constraints. Although CN can carry out traditional symbolic computation on-line, such as solving Constraint Satisfaction Problems and path planning, notice that much of the symbolic reasoning and theorem-proving may be outside the agent, in the mind of the designer. GOFAIR does not make this distinction, assuming that such symbolic reasoning occurs explicitly in, and only in, the mind of the agent.

5 Situated Perception

Whether or not Situated Agents in general, or Constraint Nets in particular, emerge as the focus of the next paradigm, the choice of target problem domain is key for moving beyond GOFAIR. 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. It should allow for situated perception, that is, perception in a specific environmental context of the relevant environmental variables. Given all that, it should also be as simple, and exciting, as possible.

One 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. By incorporating a 3D model of a telerobot's manipulator we used model-based techniques to determine the joint angles of the manipulator. It offers a cheap, fast and reliable solution to the problem of joint angle feedback [28]. Related work on visual feedback for robotics has been successful for highly constrained tasks such as table tennis [29] and throwing and juggling a ball [30] or requires special marks on the arm, special sensors or special lighting [31]. We now have a prototype system that can monitor the joint angles of the boom, stick and bucket of an excavator. We have completed a redesign, and a second prototype implementation, for a system with real-time performance at 10 Hz using parallel and distributed algorithms on image analysis boards and a Transputer system.

As the robot moves its limbs the perceptual system uses visual and proprioceptive information to provide updates to its internal self-model. A GOFAIR blocks world 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 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 off-line planning method for robot action. It is consistent with our ideas on distributed robotic architectures in Constraint Nets. So this is a truly situated robot: situated in the spatial context of its own body.

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.

This approach achieves the necessary tight coupling of perception, reasoning and action. 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.

To investigate another world in which the OFTA do not hold, Dinesh Pai and I have started the Dynamo (Dynamics and Mobile Robots) Project in our laboratory. We are experimenting with multiple mobile robots under visual control. The basic Dynamo testbed consists of fleets of radio-controlled vehicles that receive commands from a remote computer. Using a parallel and distributed SIMD/MIMD integrated environment, vision programs are able to monitor the position and orientation of each robot at 60 Hz; planning and control programs can generate and send motor commands out at 50 Hz. This approach allows umbilical-free behaviour and very rapid, lightweight fully autonomous robots. As far as we know, it is a unique and successful approach to all the tradeoffs involved



Figure 1. Two soccer players compete in the Dynamo project.
The striker on the right is shooting at the goal on the left.

in mobile robot design. In a related project we also plan to mount sensors, including television cameras, on-board the robots and transmit the data back to off-board computers. As with other experiments in mobile robotics, such as [32,33], our aim is to integrate theory and practice, as well as symbolic reasoning and control algorithms. So in a real sense these robots can see themselves and their environment, so they can monitor the effects of their own actions and the actions of others.

A long term goal is to have teams of robots engaged in cooperative and competitive behaviour. In particular, we have chosen soccer playing as one of the tasks. Our initial experiments have been successful. With Rod Barman, Stewart Kingdon, Michael Sahota and Ying Zhang, we have developed and tested path planning and motion control algorithms that allow a player to get to the ball and to shoot it at the goal, while a goalie tries to stop it, as shown in Figure 1. Some of this work is based on the Constraint Net formulation outlined above. That formulation is particularly useful here since we have written a simulation of the dynamics of the player as a constraint net and developed planning and control algorithms in CN. The Dynamo testbed will force us to develop and experiment with algorithms at all behavioral levels. Current work in the field typically adopts a hybrid scheme, grafting symbolic AI algorithms onto numerical, or fuzzy, control schemes with the problems resulting from the underlying off-line/on-line computational mismatch described earlier. We intend further practical and theoretical development of CN as a language for writing robot programs in this environment. An important hypothesis to be tested is that this single uniform on-line framework is adequate for expressing plans at all levels.

6 Conclusions

We have looked at robots looking at the world, at other robots and at themselves. We have also looked inside robots to examine their architecture and embedded assumptions. GOFAIR robots, based on the Omniscient Fortune Teller Assumptions and hybrid off-line/on-line computational models, are being challenged by Situated Agents, embedded in time and space. The Constraint Net approach models the robot and its world symmetrically as coupled dynamical systems. CN is an appropriate formalism for the new paradigm since it allows analysis of the interaction of the robot embedded in its specific world; moreover, it allows us to develop practical tools based on a unitary on-line distributed computational framework. Two systems for situated perception were described as benchmark challenges for the new approach to seeing robots.

Acknowledgments

I am grateful to Rod Barman, Craig Boutilier, Randy Goebel, Stewart Kingdon, Jim Little, Keiji Kanazawa, Peter Lawrence, David Lowe, Jane Mulligan, Dinesh Pai, David Poole, Ray Reiter, Michael Sahota, Bob Woodham and Ying Zhang for help with this. This work is supported, in part, by the Canadian Institute for Advanced Research, the Natural Sciences and Engineering Research Council of Canada and the Institute for Robotics and Intelligent Systems Network of Centres of Excellence.

References

- [1] M. B. Clowes, "On seeing things," *Artificial Intelligence*, vol. 2, pp. 79–116, 1971.
- [2] A. K. Mackworth, "On seeing things again," in *Proc. 8th International Joint Conf. on Artificial Intelligence*, (Karlsruhe, West Germany), pp. 1187–1191, 1983.
- [3] J. Haugeland, *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press, 1985.
- [4] A. K. Mackworth, "How to see a simple world: an exegesis of some computer programs for scene analysis," in *Machine Intelligence 8* (E. W. Elcock and D. Michie, eds.), pp. 510–540, New York, NY: John Wiley & Sons, 1977.
- [5] D. L. Poole, A. K. Mackworth, and R. G. Goebel, *Computational Intelligence: A Logical Approach*. Vancouver, B.C.: Dept. of Computer Science, University of British Columbia, 1992. (383 pp.).
- [6] T. M. Mitchell, "Becoming increasingly reactive," in *AAAI-90*, (Boston, MA), pp. 1051–1058, 1990.

- [7] T. S. Kuhn, *The structure of scientific revolutions*. Chicago 60637: University of Chicago Press, 1962.
- [8] R. A. Brooks, *A robot that walks: emergent behaviors from a carefully evolved network*. Cambridge, MA: Massachusetts Institute of Technology, 1988.
- [9] R. A. Brooks, "A robust layered control system for a mobile robot," *IEEE Transactions on Robotics and Automation*, vol. 2, pp. 14–23, 1987.
- [10] A. K. Mackworth, "Building robots," in *Proc. Vision Interface '92*, (Vancouver, BC), pp. 187–188, Canadian Information Processing Society, May 1992. Invited.
- [11] Y. Zhang and A. K. Mackworth, "Constraint nets: A semantic model for real-time embedded systems," Tech. Rep. TR 92-10, UBC, Vancouver, B.C., May 1992.
- [12] R. A. Brooks, "Intelligence without reason," in *IJCAI-91*, (Sydney, Australia), pp. 569–595, Aug. 1991.
- [13] R. A. Brooks, "Intelligence without representation," *Artificial Intelligence*, vol. 47, pp. 139–160, 1991.
- [14] R. Reiter and A. K. Mackworth, "A logical framework for depiction and image interpretation," *Artificial Intelligence*, vol. 41, pp. 125–155, 1990.
- [15] A. K. Mackworth, "The logic of constraint satisfaction," *Artificial Intelligence*, vol. 58, pp. 3–20, 1992.
- [16] S. J. Rosenschein and L. P. Kaelbling, "The synthesis of machines with provable epistemic properties," in *Proc. Conf. on Theoretical Aspects of Reasoning about Knowledge* (Joseph Halpern, ed.), pp. 83–98, Los Altos, CA: Morgan Kaufmann, 1986.
- [17] L. P. Kaelbling and S. J. Rosenschein, "Action and planning in embedded agents," in *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back* (P. Maes, ed.), pp. 35–48, Cambridge, MA: MIT Press, 1990.
- [18] P. E. Agre and D. Chapman, "Pengi: An implementation of a theory of activity," in *AAAI-87*, (Seattle, WA), pp. 268–272, 1987.
- [19] D. Chapman, "Vision instruction and action," Tech. Rep. MIT AI TR-1085, MIT, Cambridge, MA, June 1990.
- [20] B. C. Smith, "The owl and the electric encyclopedia," *Artificial Intelligence*, vol. 47, pp. 251–288, 1991.

- [21] D. H. Ballard, “Reference frames for active vision,” in *Proceedings IJCAI-89*, (Detroit, MI), pp. 1635–1641, 1989.
- [22] T. Winograd and F. Flores, *Understanding Computers and Cognition*. Reading, MA: Addison-Wesley, 1986.
- [23] J. Lavignon and Y. Shoham, “Temporal automata,” Tech. Rep. STAN-CS-90-1325, Stanford University, Stanford, CA, 1990.
- [24] Y. Zhang and A. K. Mackworth, “Will the robot do the right thing?,” Tech. Rep. TR 92-31, UBC, Vancouver, B.C., Nov. 1992.
- [25] I. D. Horswill and R. A. Brooks, “Situated vision in a dynamic world: Chasing objects,” in *AAAI-88*, (St. Paul, MN), pp. 796–800, 1988.
- [26] Y. Zhang and A. K. Mackworth, “Modeling behavioral dynamics in discrete robotic systems with logical concurrent objects,” in *Robotics and Flexible Manufacturing Systems* (S. G. Tzafestas and J. C. Gentina, eds.), pp. 187–196, Elsevier Science Publishers B.V., 1992.
- [27] Y. Zhang and A. K. Mackworth, “Design and analysis of embedded real-time systems: An elevator case study,” Tech. Rep. TR 93-4, UBC, Vancouver, B.C., Feb. 1993.
- [28] I. J. Mulligan, A. K. Mackworth, and P. D. Lawrence, “A model-based vision system for manipulator position sensing,” in *Proc. IEEE Workshop on Interpretation of 3D Scenes*, (Austin, TX), pp. 186–193, 1989.
- [29] R. L. Andersson, *A Robot Ping-Pong Player: Experiment in Real-Time Intelligent Control*. Cambridge, MA: MIT Press, 1988.
- [30] E. W. Aboaf, A. K. Drucker, and C. B. Atkeson, “Task-level robot learning: Juggling a tennis ball more accurately,” in *Proc. IEEE Int. Conf. on Robotics and Automation*, pp. 1290–1295, 1989.
- [31] J. M. Hollerbach, “A review of kinematic calibration,” in *The Robotics Review 1* (O. Khatib and J.J. Craig and T. Lozano-Perez, ed.), pp. 207–242, Cambridge, MA: MIT Press, 1989.
- [32] O. Amidi and C. Thorpe, “Integrated mobile robot control,” in *SPIE Mobile Robots V*, pp. 504–523, 1990.
- [33] T. Skewis and V. Lumelsky, “Experiments with a mobile robot operating in a cluttered unknown environment,” in *Proc. 1992 IEEE Int. Conf. on Robotics and Automation*, (Nice, France), pp. 1482–1487, May 1992.