

Agents, Bodies, Constraints, Dynamics, and Evolution

Alan K. Mackworth

■ *The theme of this article is the dynamics of evolution of agents. That theme is applied to the evolution of constraint satisfaction, of agents themselves, of our models of agents, of artificial intelligence, and, finally, of the Association for the Advancement of Artificial Intelligence (AAAI). The overall thesis is that constraint satisfaction is central to proactive and responsive intelligent behavior.*

It was a great pleasure to serve as president of the Association for the Advancement of Artificial Intelligence (AAAI). I think that everyone who serves begins with great trepidation but leaves with a sense of satisfaction. I certainly did. One of the sources of satisfaction was the opportunity to give the presidential address at AAAI-07 in Vancouver, my hometown. This article is based on that talk. That opportunity allowed me to step back and think of the big picture, the long perspective. The way to read the title, “Agents, Bodies, Constraints, Dynamics, and Evolution,” as it is intended to be read, is to say that the article is about *agents with bodies subject to constraints on their dynamics as they undergo evolution*. Thus, I am telling you a story. It is a story that does not, I hope, tell any lies by omission or commission. It is idiosyncratic, that being the nature of narrative, to make it fit what I want to say as I make sense of that big picture. I hardly need add that this is a highly personal and telegraphic view of the topics I discuss. In a sense, it is a cartoon.

The theme of this article, then, is the dynamics of evolution. I shall apply that theme to five different topics: to the evolution of the idea of constraint satisfaction, to agents themselves as they evolve, to our own models of agents and how they have changed over time, to the field of AI, and, finally, I shall apply it to AAAI, the organization. Thus, we can imagine each of these concepts, including the field of AI and AAAI, as agents acting in an environment. An agent in an active environment changes that environment and, in turn, is changed by that environment as each evolves in time. To tie it all together, I do have a thesis, and it will not surprise too many of you, perhaps, when I say that the thesis is *constraint satisfaction is central to intelligent behavior*.

Constraint Satisfaction Problems

What may surprise you, however, is that constraint satisfaction of today is not constraint satisfaction as it was 30 years ago—what we

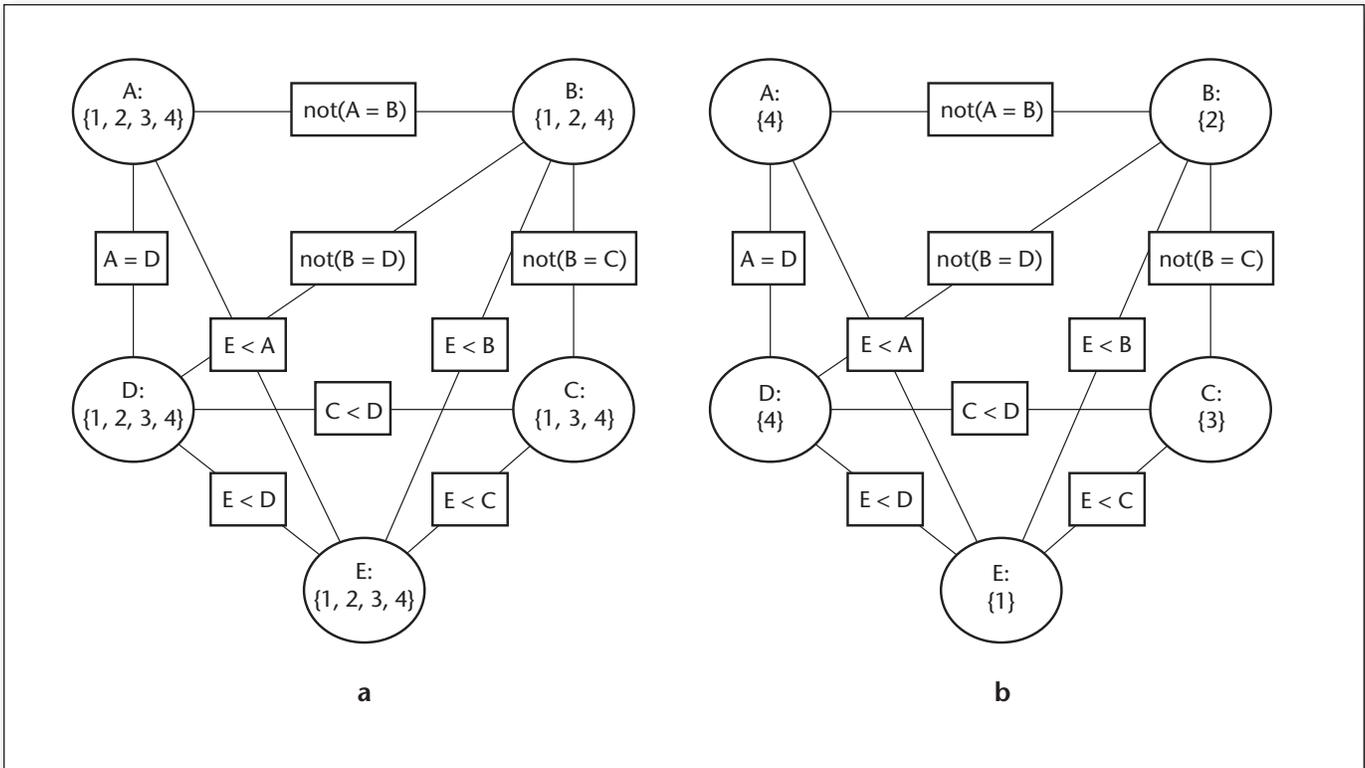


Figure 1. A Meeting Scheduling Problem.

(a) As a CSP in AIspace. (b) The solution to the scheduling problem found with arc consistency.

might call *good old-fashioned constraint satisfaction* (GOFCS). Constraint satisfaction itself has evolved far beyond GOFCS. However, I will start with GOFCS as exemplified in the constraint satisfaction problem (CSP) paradigm. The whole concept of constraint satisfaction is a powerful idea. It arose in several fields somewhat simultaneously; a group of us in the early 1970s abstracted the underlying model. Simply, many significant sets of problems of interest in artificial intelligence can each be characterized as a CSP with a set of variables; each variable has a domain of possible values, and there are various constraints on those variables, specifying which combinations of values for the variables are allowed (Mackworth 1977). The constraints may be between two variables or they may be among more than two variables. Consider a simple scheduling CSP example with five one-hour meetings *A*, *B*, *C*, *D*, and *E* to be scheduled. There are constraints such as meeting *A* must occur after meeting *E*; meetings *A* and *D* have to occur at the same time but meetings *A* and *B* cannot occur at the same time, and so on. A CSP can be represented as a bipartite graph as shown for the scheduling problem in figure 1(a), where the meetings (variables) are shown as elliptical nodes and the constraints are shown as rectangular nodes.

As domains of possible values for the meeting

start times we have one, two, three, or four o'clock in the afternoon. We have to find a solution, a set of times that satisfies the constraints for each meeting. If you are interested, you can go to our AIspace website¹ and use the consistency-based CSP solver applet we have developed. There you can see this problem satisfied using arc consistency, and other techniques, to find out whether there is no solution, a unique solution, or more than one. For this problem, as it happens, there is a unique solution for the set of possible meetings, found by running arc consistency, shown in figure 1(b).

Using the applet, one can see the constraints propagate around, based on arc consistency. Our design of this AIspace applet for CSPs was inspired by a historical video that David Waltz made in the early 1970s to animate his edge-labeling algorithm for line drawings of blocks world scenes (Waltz 1975). That video shows the constraints propagating around a line drawing as the sets of possible corners at each vertex are pruned by the edge-label constraints.

Here is a familiar CSP example from current popular culture. A typical Sudoku puzzle is shown in figure 2.

The solver has to fill in the squares, each with a digit chosen from {1, 2, ..., 9}, where the constraints are that every row, every column, and

every 3 x 3 subgroup has to be a permutation of those nine digits.

Usually, to avoid torturing the puzzle solver, the puzzle maker ensures there is one and only one solution to the problem. One can find these solutions using arc consistency techniques and search; moreover, one can easily generate and test potential Sudoku puzzles to make sure they have one and exactly one solution before they are published. AI has its uses.

Arc consistency is a simple member of the class of algorithms we called *network consistency algorithms*. The basic idea is that one can, before constructing global solutions, efficiently eliminate local nonsolutions. The constraints are all conjunctive, so each value for every variable must satisfy the constraints. And since all of the constraints have to be satisfied, if there is any local value configuration that does not satisfy them, one can throw that tuple out; that is a “no good.” So one can discover (learn) those local inconsistencies very quickly, in linear, quadratic, or cubic time. Those savings will give huge, essentially exponential savings when one does start searching, constructing global solutions, using backtracking or other approaches. The simplest algorithm is arc consistency, then path consistency, then *k*-consistency, and so on. Many other AI researchers contributed to this development, including Richard Fikes, Dave Waltz, Ugo Montanari, and Eugene Freuder. For a detailed historical perspective on that development see Freuder and Mackworth (2006). Since those early days network consistency algorithms have become a major research industry. The CSP industry is one of the largest in AI, and in many conferences over the years it has been the single largest topic with the greatest number of sessions and papers. In fact, it has now evolved into its own field of computer science and operations research, called *constraint programming*. The CSP approach has been combined with logic programming and various other forms of constraint programming. Indeed, it has come to have its own journals, its own conferences, and all the other accoutrements of a fully fledged academic field. It is having a major impact in many industrial applications of AI, logistics, planning and scheduling, and combinatorial optimization. For a comprehensive overview, see the *Handbook of Constraint Programming* (Rossi, van Beek, and Walsh 2006). However, this discussion, although pertinent to insight into the development of AI, is not central to my theme here.

Pure Good Old-Fashioned AI and Robotics (GOFAIR)

Here, we are more interested in how we build agents and how the way we build them has

5	3			7				
6			1	9	5			
	9	8					6	
8				6				3
4			8		3			1
7				2				6
	6					2	8	
			4	1	9			5
				8			7	9

Figure 2. A Sudoku Puzzle.

evolved over time. John Haugeland (1985) was the first to use the phrase *good old-fashioned AI* (GOFAI) in talking about symbolic AI using reasoning and so on as a major departure from early work in cybernetics and control theory. GOFAI has come to be a straw man for advocates of subsymbolic approaches. AI at that point, when we discovered these symbolic techniques, tended to segregate itself from those other areas. Lately, however, we have been coming back together in some ways, and that is another theme of this article. Let me quickly add here that there was a great deal of early work in symbolic programming of robots—a lot of great work. That work can be characterized as good old-fashioned AI and robotics (GOFAIR) (Mackworth 1993). The way I shall describe it here is just a cartoon, for as I said, I am telling a story here not meant to be taken literally.

GOFAIR Metaassumptions

In a cartoon sense, a pure GOFAIR robot operates in a world that satisfies several metaassumptions: (1) there is a single agent; (2) the agent executes actions serially; (3) actions occur in a deterministic world; (4) the world is a fully observable, closed world; (5) the agent has a perfect internal model of infallible actions and world dynamics; (6) perception is needed only to determine the initial world state; and (7) a perfect plan to achieve the goal is obtained by reasoning, and executed in an open loop.

There is a single agent in the world that executes its actions serially. It does not have two hands that can work cooperatively. The world is deterministic.

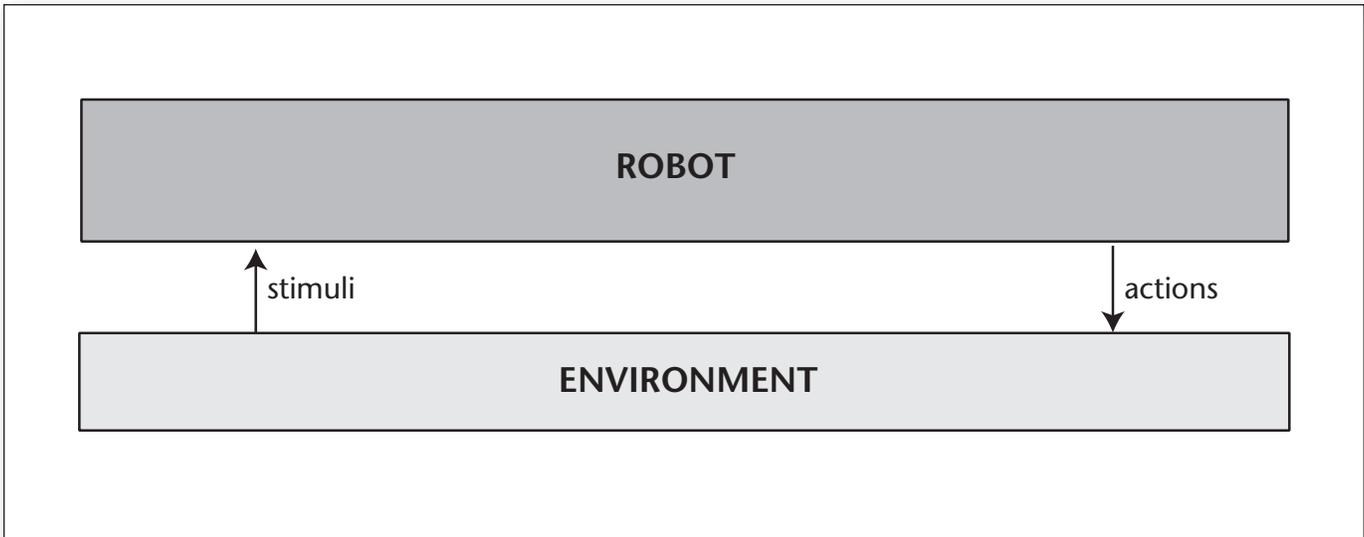


Figure 3. A Robot Coevolving with Its Environment

It is fully observable. It is closed, so if I do not know something to be true, then it is false, thanks to the “closed world assumption” (Reiter 1978). The agent itself has a perfect internal model of its own infallible actions and the world dynamics, which are deterministic. If these assumptions are true then perception is needed only to determine the initial world state. The robot takes a snapshot of the world. It formulates its world model. It reasons in that model and it can combine that reasoning with its goals using, say, a first order theorem-prover to construct a plan. This plan will be perfect because it will achieve a goal even if it executes the plan open loop. So, with its eyes closed, it can just do action *A*, then *B*, then *C*, then *D*, then *E*. If it happens to open its eyes again, it would see, “Oh, I did achieve my goal, great!” However, there is no need for it to open its eyes because it had a perfect internal model of these actions that have been performed, and they are deterministic, and so the plan was guaranteed to succeed with no feedback from the world.

CSPs and GOFAIR

What I would like you, the reader, to do is to think of the CSP model as a very simple example of GOFAIR. There are no robots involved, but there are some actions. The solver is placing numbers in the squares and so on. In pure GOFAIR there is a perfect model of the world and its dynamics in the agent’s head, so I call the agent then *an omniscient fortune teller*, as it knows all and it can see the entire future because it can control it, perfectly. Therefore if these conditions are all satisfied, then the agent’s world model and the world itself will be in perfect

correspondence—a happy state of affairs, but it doesn’t usually obtain. However, when working in this paradigm we often failed to distinguish the agent’s world model and the world itself because there really is no distinction in GOFAIR. We confused the agent’s world model and the world—a classic mistake.

A Robot in the World

Now we come to think about the nature of robots. A robot acts in a world. It changes that world, and that world changes the robot. We have to conceive of a robot in an environment, performing actions in an environment, and the environmental stimuli, which could be sensory stimuli or physical stimuli, will change the robot. Therefore, think of the robot and its environment as two coupled dynamic systems, operating in time, embedded in time, each changing the other, as they coevolve as shown in figure 3. They are mutually evolving, perpetually, to some future state, because, of course, the environment could contain many other agents who see this robot as part of their environment.

Classic Horizontal Architecture

Again, in a cartoon fashion, consider the so-called three boxes model or the horizontal architecture model for robots. Since perception, reasoning, and action are the essential activities of any robot—why not just have a module for each?

As shown in figure 4, the perception module interprets the stimuli coming in from the environment; it produces a perfect three-dimensional world model that is transmitted to the reasoning module, which has goals, either internally gener-

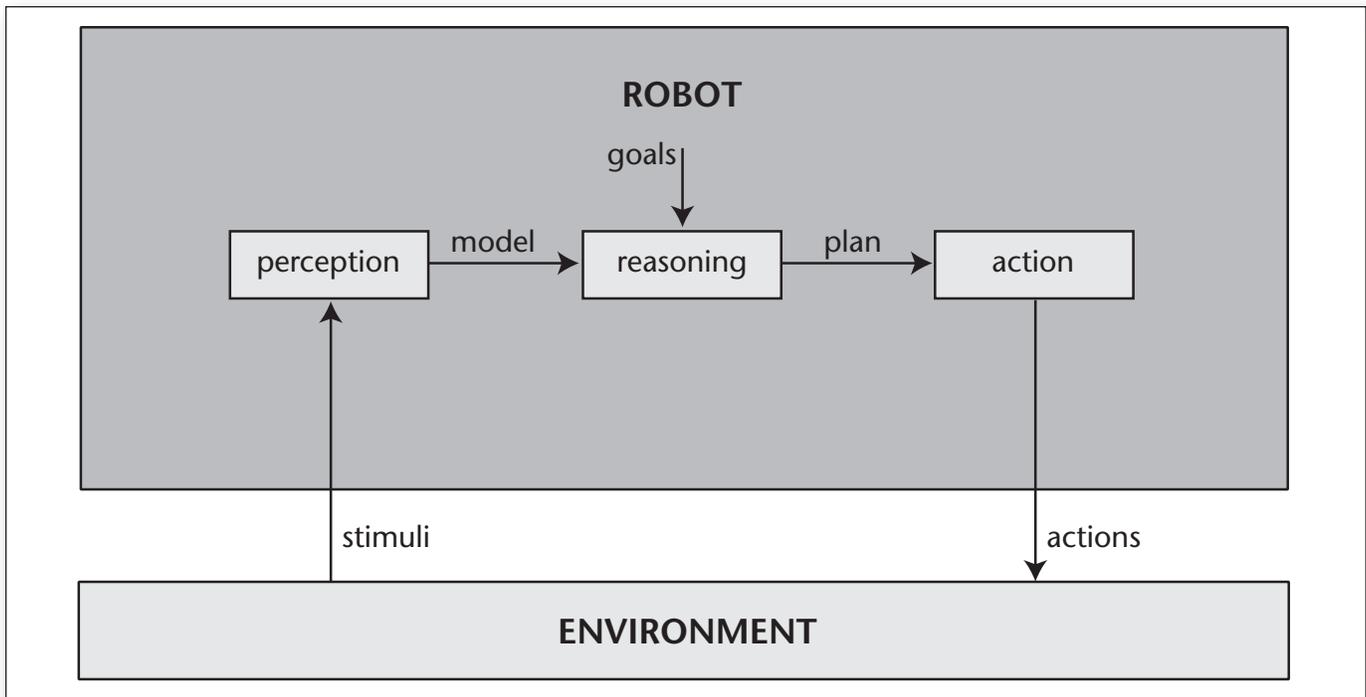


Figure 4. A Horizontal Architecture for a GOFAIR Robot.

ated or from outside. Combining the model and the goals, the reasoning module produces a plan. Again, that plan is just a sequence of the form: Do this, do this, do this, then stop. There are no conditionals, no loops in these straight-line plans. Those actions will, when executed, change the world perfectly according to the goals of the robot. Now, unfortunately for the early hopes for this paradigm, this architecture can only be thought of as a really good first cut. You know that if you wanted to build a robot it is a really good first thought. You want to push it as hard as you can, because it is nice and simple, it keeps it clean and modular, and all the rest of it. It is simple but, unfortunately, not adequate. Dissatisfaction with this approach drove the next stage of evolution of our views of robotic agents, illustrating our theme.

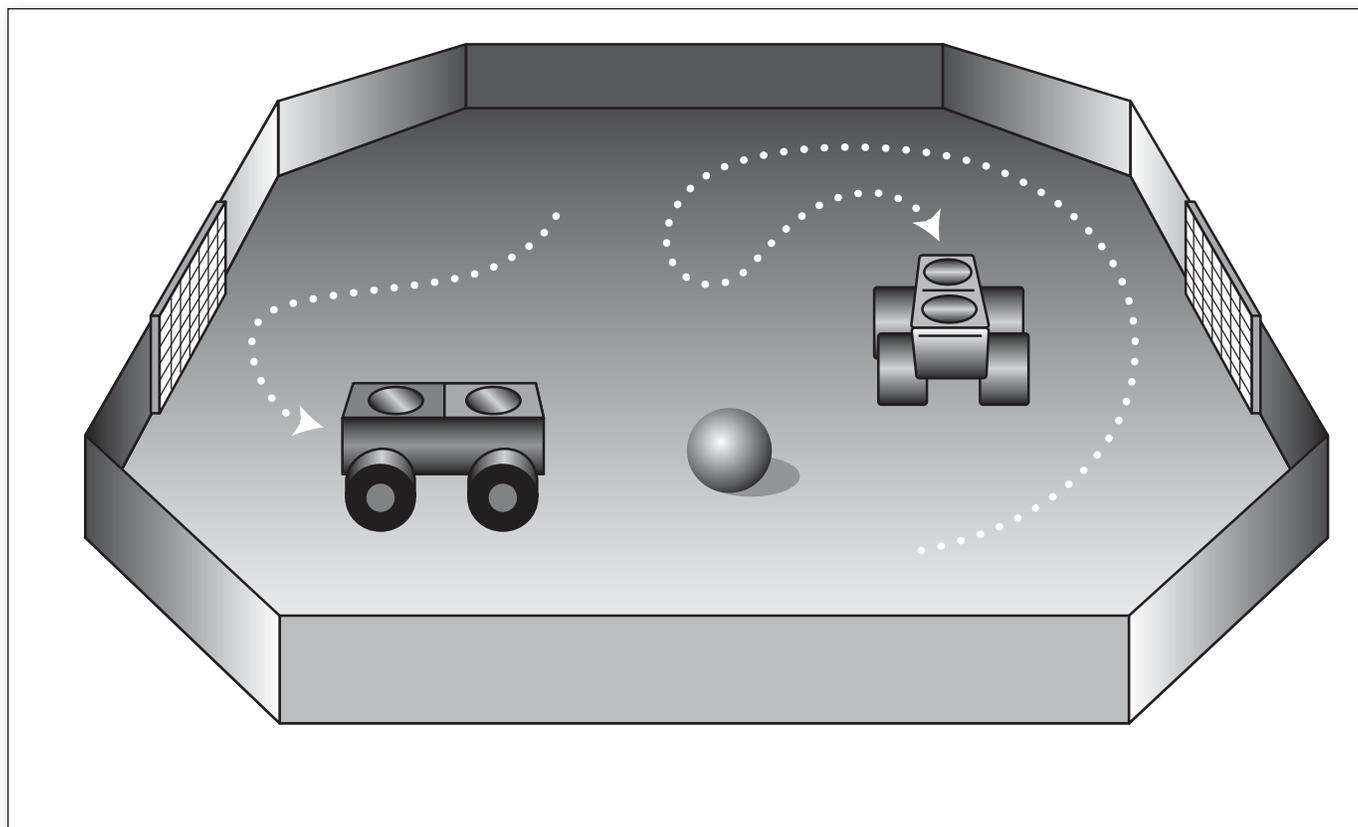
The Demise of GOFAIR

GOFAIR robots succeed in controlled environments such as blocks worlds and factories, but they cannot play soccer! GOFAIR did work, and does work, as long as the blocks are matte blocks with very sharp edges on black velvet backgrounds. It works in factories if there is only one robot arm and it knows exactly where things are and exactly where they are going to go. The major defect, from my point of view, is that GOFAIR robots certainly cannot, and certainly never will, play soccer. I would not let them into my home without adult

supervision. In fact, I would advise you not to let them into your home either.

It turns out that John Lennon, in retrospect, was a great AI researcher since in one of his songs he mused, "Life is what happens to you when you're busy making other plans." (Lennon 1980. The key to the initial success of GOFAIR is that the field attacked the planning problem and came up with really powerful ideas, such as GPS, STRIPS, and back-chaining. This was revolutionary. Algorithms were now available that could make plans in a way we could not do before. Miller and colleagues' book *Plans and the Structure of Behaviour* (Miller, Galantner, and Pribam 1960) was a great inspiration and motivation for this work. In psychology there were few ideas about how planning could be done until AI showed the way. The GOFAIR paradigm demonstrated how to build proactive agents for the very first time.

But planning alone does not go nearly far enough. Clearly, a proactive GOFAIR robot is indeed an agent that can construct plans and act in the world to achieve its goals, whether short term or long term. Those goals may be prioritized. That is well and good. However, "There are more things in heaven and earth, Horatio, than are dreamt of in your philosophy."² In other words, events will occur in the world that an agent does not expect. It has to be able to react quickly to interrupts from the environment, to real-time changes, to immi-



Drawing by Jim Marin, ©2009.

Figure 5. The Setup for the World's First Robot Soccer Game, between Zeno and Heraclitus.

nent threats to safety to itself or humans, to other agents, and so on. An intelligent robot must be both proactive *and* responsive. An agent is proactive if it acts to construct and execute short-term and long-term plans and achieve goals in priority order. An agent is responsive if it reacts in real time to changes in the environment, threats to safety, and other agents' actions.

Beyond GOFAIR to Soccer

So that was the real challenge that was before us in the 1980s to the GOFAIR cartoon worldview. How could we integrate proactivity and reactivity? In 1992, I made a proposal (Mackworth 1993) that it is fine to say robots must be proactive and reactive or responsive but we needed a simple task domain in order to force us to deal with those kinds of issues. I proposed robot soccer as that domain in that paper. Actually, I proposed it after we had actually already built soccer players in our lab and made them work. We built the world's first robot soccer players, using cheap toy radio-controlled monster trucks, as drawn in figure 5. The first two players were named after Zeno and Heraclitus. You can see videos of the first robot soccer games on the web.³

A single color camera, looking down on these trucks, could see the colored circles on top of the trucks so that the perceptual system could distinguish Zeno from Heraclitus. It could also see the ball and the goals. Each truck has its own controller. Since they cannot turn in place—they are nonholonomic—it is very hard actually to control them. It is a very tricky problem to control this kind of steerable robot. The path-planning problems have to be solved in real time. Of course, one is trying to solve a path-planning problem, as the ball is moving and the opponent is moving in order to get that ball; that is very tricky computationally. We were pushing the limits both of our signal-processing hardware and the CPUs to get this to work in real time. We were running there at about 15 Hz cycle time. The other problem was that our lab was not big enough for these monster trucks. We wanted to go with more trucks, but the department head would not give us a bigger lab. So we were forced to go to smaller robots, namely, 1/24th scale radio-controlled model Porsches, that we called *Dynamites*. These cars were running on a ping-pong table with a little squash ball. In the online video, one can see the players alternating between offensive and defensive behaviors. The

behaviors the robots exhibit are clearly a mix of proactive and responsive behaviors, demonstrating the theme of evolution of our models of agents beyond the GOFAIR approach.

Incidentally, there was an amazing, successful effort to get chess programs to the point where they beat the world champion (Hsu 2002). But from the perspective presented here, it changes only the single-agent Sudoku puzzle into a two-agent game; however, all the other aspects of the Sudoku domain remain the same—perfect information, determinism, and the like. Chess loses its appeal as a domain for driving AI research in new directions.

We managed to push all our soccer system hardware to the limit, so we were able to develop two-on-two soccer. The cars were moving at up to 1 meter per second and autonomously controlled at 30 Hz. Each had a separate controller off board, and the cars were entirely independent. The only thing they shared is a common front-end vision-perceptual module. We were using transputers (a 1 MIP CPU) because we needed significant parallelism here. You can see a typical game segment with the small cars on the web.⁴ We were able to do the real-time path planning and correction and control at about 15–30 Hz, depending, but that was really the limit of where we could go at that time (1992–94) because we were limited by the hardware constraints.

RoboCup

As it happens, shortly thereafter, some Japanese researchers started to think along similar lines. They saw our work and said, “Looks good.” Instead of using steerable robots, as we had, they chose holonomic robots that can spin in place. Hiroaki Kitano and his colleagues in Japan proposed RoboCup (Kitano 1998). In Korea, the MiroSot group⁵ was also intrigued by similar issues. It made for an interesting international challenge.

The first RoboCup tournament was held in Nagoya, Japan, in 1997. Our University of British Columbia (UBC) team participated; it was a milestone event. Many researchers have subsequently made very distinguished contributions in the robot soccer area, including Peter Stone, Manuela Veloso, Tucker Balch, Michael Bowling, Milind Tambe, and many others. It has been fantastic. At RoboCup 2007, held in Atlanta, Georgia, there were approximately 2,700 participant agents, and of those, about 1,700 were people and 1,000 were robots. A review of the first 10 years of RoboCup has recently appeared (Visser and Burckhard 2007) showing how it has grown in popularity and influenced basic research.

RoboCup has become incredibly exciting, a little cutthroat and competitive with, perhaps, some dubious tactics at times but that is the nature of

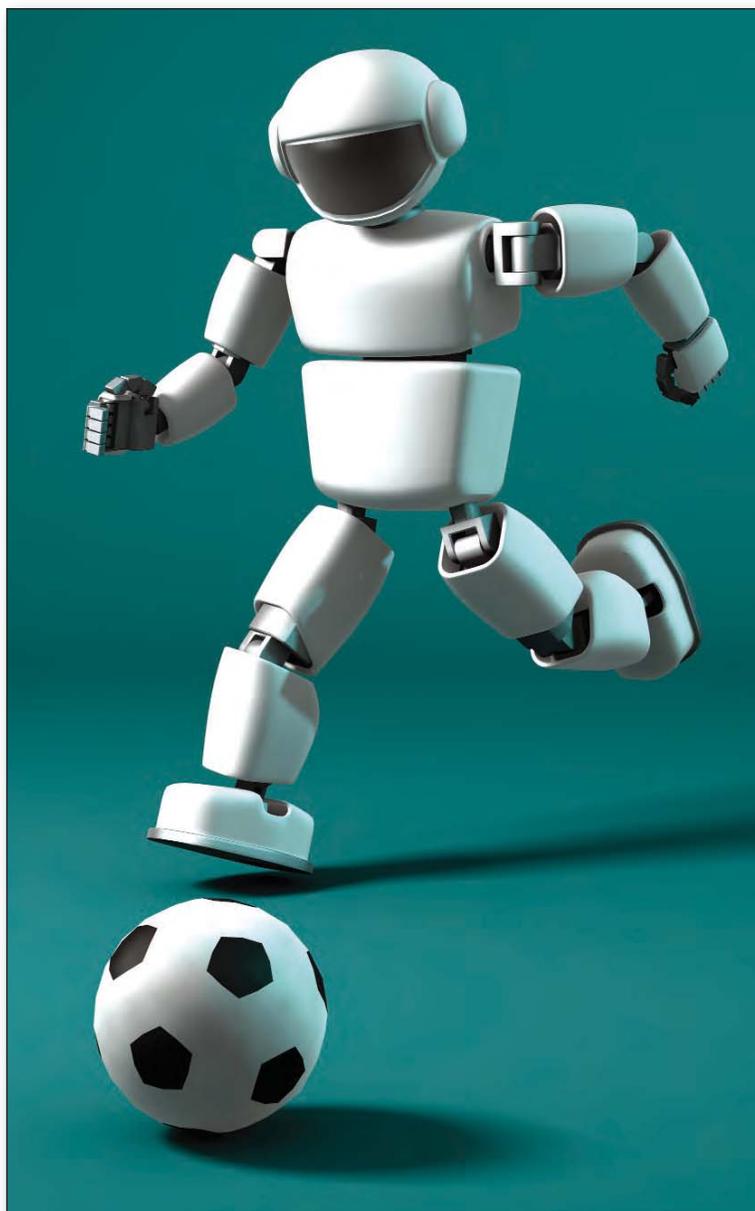


Illustration by Giacomo Marchesi, ©2009.

Figure 6. Humanoid Robot Soccer Player.

intense competition in war and soccer. More importantly, robot soccer has been incredibly stimulating to many young researchers, and it has brought many people into the field to do fine work including new competitions such as RoboRescue and RoboCup@Home. The RoboCup mission is to field a team of humanoid robots to challenge and beat the human champions by 2050, as suggested by figure 6.

From Sudoku to Soccer and Beyond

Now let us step back a bit and consider our theme of evolutionary development. If one thinks of the

	Sudoku	Soccer
Number of agents	1	23
Competition	No	Yes
Collaboration	No	Yes
Real time	No	Yes
Dynamics	Minimal	Yes
Chance	No	Yes
Online	No	Yes
Planning Horizons	No	Yes
Situated Perception	No	Yes
Partially Observable	No	Yes
Open World	No	Yes
Learning	Some	Yes

Figure 7. Comparison of Sudoku and Soccer along Various Dimensions.

Sudoku domain as the exemplar of GOFAIR in a very simple-minded way, then soccer is an exemplar of something else—what is the something else? I think it is situated agents, and so we are in transition from one paradigm to another. As shown in figure 7, we can compare Sudoku and soccer as exemplar tasks for each paradigm, GOFAIR and situated agents, respectively, along various dimensions.

I shall not go through these dimensions exhaustively. In soccer we have 23 agents: 22 players and a referee. Soccer is hugely competitive between the teams obviously, but also of major importance is the collaboration within the teams, the teamwork being developed, the development of plays, and the communications systems, the signaling systems between players and the protocols for them. Soccer is very real time. There is a major influence of dynamics and of chance. Soccer is online in the sense that one cannot compute a plan offline and then execute it, as one can in GOFAIR. Whenever anything is done, a plan almost always must be recomputed. There exist a variety of temporal planning horizons, from “Can I get my foot to the

ball?” through to “Can I get the ball into the net?” and “Can I win this tournament?” The visual perception is very situated and embodied. Vision is on board the robots now in most of the leagues, so a robot sees only what is visible from where it is, meaning the world is obviously only partially observable. The knowledge base is completely open because one cannot infer much about what is going on behind one’s back. The opportunities for robot learning are tremendous.

From GOFAIR to Situated Agents

As we make this transition from GOFAIR to situated agents, how is it done? There has been a whole community working on situated agents, since James Clerk Maxwell in the late 19th century building governors for steam engines and the like. Looking at Maxwell’s classic paper “On Governors” (Maxwell 1868), it is clear that he produced the first theory of control, trying as he was to understand why James Watts’s feedback controller for steam engines actually worked, under what conditions it was stable, and so on. Control theorists have had a great deal to say about situated agents for the last century or so. So, one way to build a situated agent would be to suggest that we put AI and control together: to stick a planner, an AI planner, GOFAIR or not, on top of a reactive control-theoretic controller doing PID control. One could also put in a middle layer of finite-state mode control. These are techniques we fully understand, and that is, in fact, how we did it for the first soccer players, which I have described above. There was a two-level controller. However, there are many problems with that approach, not the least being debugging it, understanding it, let alone proving anything about it. It was all very much: try it and see. It was very unstable as new behaviors were added. It had to be restructured at the higher level and so on. Let me just say that it was a very graduate student intensive process requiring endless student programming hours! So rather than gluing a GOFAIR planner on top of a multilayer control-theoretic controller we moved in a different direction.

I argued that we must abandon the metaassumptions of GOFAIR but keep *constraint satisfaction*. My response was that we just give up on those metaassumptions of GOFAIR but not throw out the baby of constraint satisfaction with the bathwater of the rest of GOFAIR. Constraint satisfaction was, and is, the key in my mind, and the reason it is the key is that we understand symbolic constraints, as well as numerical. We understand how to manipulate them and so on. We understand even first-order logic as a constraint-solving system, thanks to work on that side, but we also understand constraints in the control world. We understand that a thermostat is trying to solve a



Photo courtesy Stefan Eissing.

Figure 8. Robot Friends Playing Soccer.

constraint. We have now a uniform language of constraint solving or satisfaction although one aspect may be continuous while the other may be discrete or even symbolic. There is a single language or single paradigm to understand it from top to bottom, which is what we need to build clean systems. The constraints now though are dynamic: coupling the agent and its environment. They are not like the timeless Sudoku constraint: every number must be different now and forever. When one is trying to kick a ball, the constraint one is trying to solve is whether the foot position is equal to the ball's position at a certain orientation, at a certain velocity, and so on. Those are the constraints one is trying to solve, and one really does not care how one arrives there. One simply knows that at a certain point in time, the ball will be at the tip of the foot, not where it is now but where it will be in the future. So this is a constraint, but it is embedded in time, and it is changing over time as one is trying to solve it, and clearly, that is the tricky part.

Thus, constraints are the key to a uniform archi-

ecture, and so we need a new theory of constraint-based agents. This has set the stage. I shall leave you in suspense for a while for a digression before I come back to sketch that theory. Its development is part of the evolutionary process that is the theme of this article.

Robot Friends and Foes

I digress here briefly to consider the social role of robots. Robots are powerful symbols; they have a really interesting emotional impact. One sees this instinctively if one has ever worked with kids and Lego robotics or the Aibo dogs that we see in figure 8, or with seniors who treat robots as friends and partners. We anthropomorphize our technology with things that look almost like us or like our pets—although not too much like us. (That is the “uncanny valley” [Mori 1982]). We relate to humanoid robots very closely emotionally. Children watching and playing with robot dogs appear

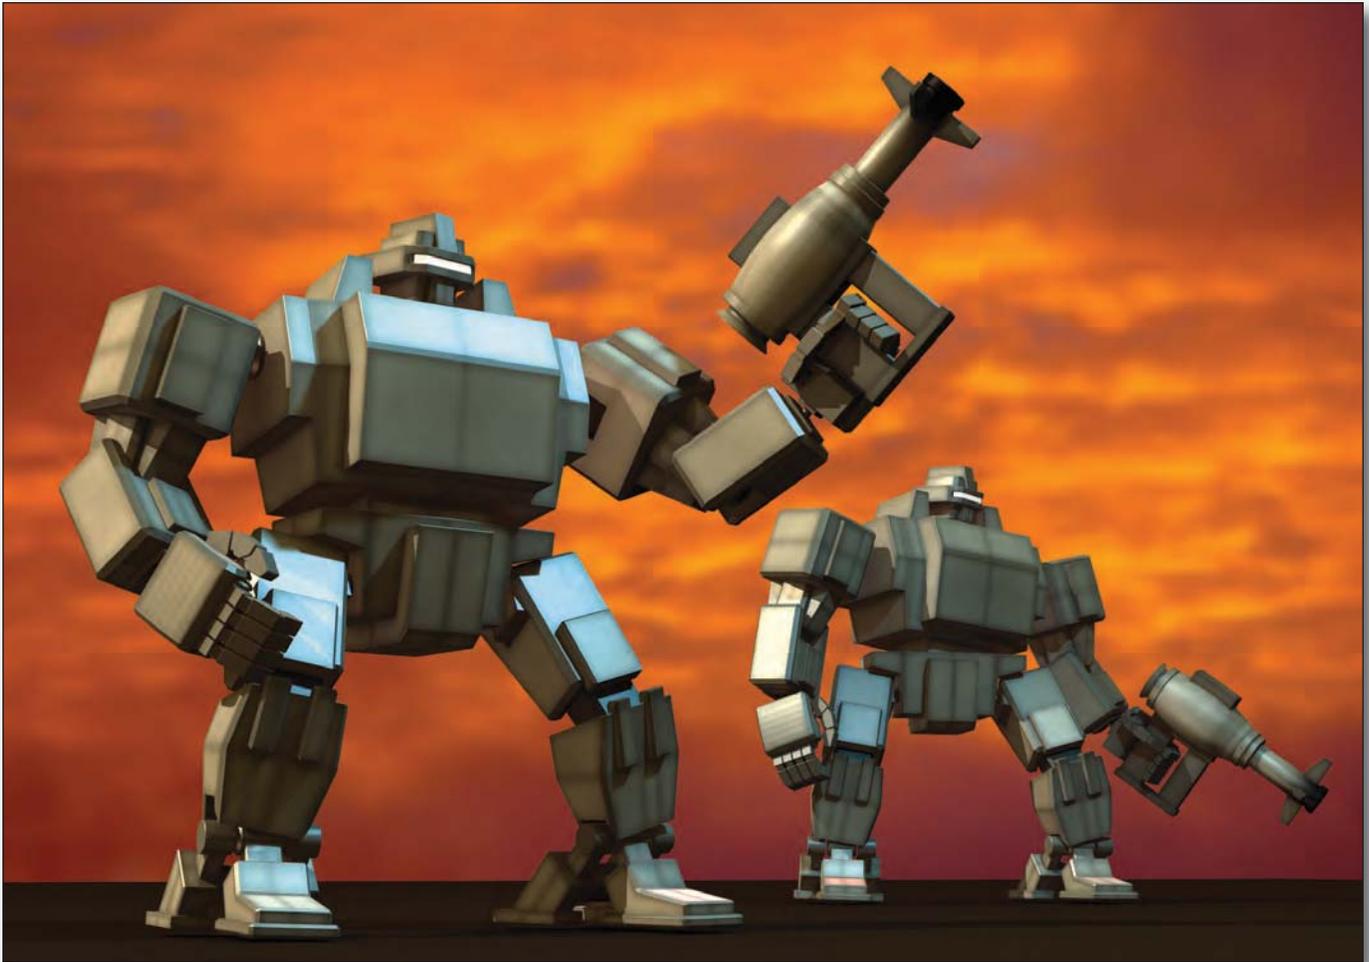


Illustration by Giacomo Marchesi, ©2009.

Figure 9. ... and Robot Foes.

to bond with them at an emotional level.

But, of course, the flip side is the robot soldier (figure 9), the robot army, and the robot tank.

Robots, Telerobots, Androids, and Cyborgs

Robots really are extensions of us. Of course, there are many kinds of robots. One uses the word *robot* loosely but, technically, one can distinguish between strictly autonomous robots and telerobots where there is human supervisory control, perhaps at a distance, to Mars or in a surgical situation, for example. There are androids that look like us and cyborgs that are partly us, partly machine. The claim is that robots are really reflections of us, and that we project our hopes and fears onto them. That this has been reflected in literature and other media over the last two centuries is a fact. I do not need to bring to mind all the robot movies, but robots do stand as symbols for our technology.

Dr. Frankenstein and his creation, in *Franken-*

stein; or, The Modern Prometheus (Shelley 1818), stood as a symbol of our fear, a sort of Faustian fear that that kind of power, that kind of projection of our own abilities in the world, would come back and attack us. Mary Shelley's work explored that, and Charlie Chaplin's *Modern Times* (1936) brought the myth up to date. Recall the scene in which Chaplin is being forced to eat in the factory where, as a factory worker, his entire pace of life is dictated by the time control in the factory. He is a slave to his own robots, and his lunch break is constrained because the machines need to be tended. He is, in turn, tended by an unthinking robot who keeps shoving food into his mouth and pouring drinks on him until finally, it runs amok. Chaplin was making a very serious point that our technology stands in real danger of alienating and repressing us if we are not careful.

I shall conclude this somewhat philosophical interjection with the observations of two students of technology and human values. Marshall



Photo Courtesy University of Calgary.

Figure 10. RoboSurgeon NeuroArm.

McLuhan argued, although he was thinking of books, advertising, television, and other issues of his time but it applies equally to robots, “We first shape the tools and thereafter our tools shape us.” (McLuhan 1964). Parenthetically, this effect can be seen as classic projection and alienation in the sense of Ludwig Andreas Feuerbach (1854).

The kinds of robots we build, the kinds of robots we decide to build, will change us as they will change our society. We have a heavy responsibility to think about this carefully. Margaret Somerville is an ethicist who argues that actually the whole species *Homo sapiens* is evolving into *Techno sapiens* as we project our abilities out (Somerville 2006). Of course, this is happening at an accelerating rate. Many of our old ethical codes are broken and do not work in this new world, whether it is in biotechnology or robotics or almost any other area of technology today. As creators of some of this technology, it is our responsibility to pay serious attention to that.

Robots: One More Insult to the Human Ego?

Another way of thinking about our fraught and ambivalent relationship with robots is that this is really one more insult. How much more can humankind take? Robotics is only the latest displacement of the human ego from center stage. Think about the intellectual lineage that links Copernicus, Darwin, Marx, Freud, and robots. This may be a stretch in thinking but perhaps not.

Humans thought they were at the center of the universe until Nicolaus Copernicus proposed that the earth was not at the center and the sun should be seen that way. Charles Darwin hypothesized we are descended from apes. Karl Marx claimed that many of our desires and goals are determined by our socioeconomic status, and, thus, we are not as free as we thought. Sigmund Freud theorized that one’s conscious thoughts are not freely chosen but rather come from the unconscious mind. Now I suggest that you can think of robots as being in



Figure 11. Prototype Smart Wheelchair (UBC, 2008).

that same great lineage of saying, you, *Homo sapiens*, are not unique. Now, there are other entities, robots, created by us, that can also perceive, think, and act. They could become as smart as we are. But this kind of projection can lead to a kind of moral panic: “The robots are coming! The robots are coming! What are we going to do?” When we talk to the media the first questions reporters ask are, typically: “Are you worried about them rising up and taking over?” and “Do you think they’ll keep us as pets?” The public perception of robots is evolving as our models of robots and the robots themselves evolve.

Helpful Robots

To calm this kind of panic we need to point to some helpful robots. The arm shown in figure 10 (on the previous page), the RoboSurgeon NeuroArm, is actually fabricated from nonmagnetic parts so it can operate within a magnetic resonance imaging (MRI) field. The surgeon is able to do neurosurgery telerobotically, getting exactly the right parts of the tumor while seeing real-time feedback as the surgery is performed. An early prototype of our UBC smart wheelchair work is shown in figure

11. This chair can use vision and other sensors to locate itself, map its environment, and allow its user to navigate safely.

RoboCars: DARPA Urban Challenge

Continuing with the helpful robot theme, consider autonomous cars. The original DARPA Challenges, in 2004 and 2005, and the Urban Challenge in 2007, have catalyzed significant progress, stimulated by Ron Brachman (my predecessor as AAI president) at DARPA. Sebastian Thrun and his team at Stanford developed Junior (figure 11a), loaded with sensors and actuators and horsepower and CPUs of all sorts, facing off against Boss (figure 11b), the Carnegie Mellon University/General Motors Tartan racing team in the fall of 2007, with Boss taking first place and Junior second in the Urban Challenge.⁶ The media look at these developments and see them as precursors to robot tanks, cargo movers, and automated warfare, reaching an understanding of why DARPA funded them. However, Thrun (2006) is an evangelist for a different view of such contests. The positive impact of having intelligent cars would be enormous. Consider the potential ecological savings of using highways so much more efficiently instead of paving over farmland. Consider the safety aspect in reducing the annual carnage of 4,000 road accident deaths a year in Canada alone. Consider the fact that cars could negotiate at intersections, the way Dresner and Stone (2008) have simulated to show you could get maybe two to three times the throughput in cities in terms of traffic if these cars could talk to each other instead of having to wait for dumb stop signs and traffic lights. Consider the ability for the elderly or disabled to get around on their own. Consider the ability to send one’s car to the parking lot by itself and then call it back later. There would be automated warehouses for cars instead of using all that surface land for parking. Truly, the strong positive implications of success in this area are enormous. But, can we trust them? This is a real problem and it is one of the major problems. In terms of smart wheelchairs, one major reason they do not already exist now is liability. It is almost impossible to get an insurance company to back a project or a product. This clarifies why the car manufacturers have moved very slowly in an incremental way to develop intelligent technology.

Can We Trust Robots?

There are some real reasons we cannot yet trust robots. The way we build them now, not only are they not trustworthy, they are also unreliable. So, can they do the right thing? Will they do the right thing? And then, of course, there is the fear that I alluded to earlier that eventually they will become autonomous, with free will, intelligence, and consciousness.



Photos courtesy US Defense Advanced Research Projects Agency.

Figure 12. Two Competitors in the DARPA Urban Challenge.

Ethics at the Robot-Human Interface

Do we need robot ethics, for us and for them? We do. Many researchers are working on this (Anderson and Anderson 2007). Indeed, many countries have suddenly realized this is an important issue. There will have to be robot law. There are already robot liability issues. There will have to be professional ethics for robot designers and engineers just as there are for engineers in all other disciplines. We will have to factor the issues around what we should do ethically in designing, building, and deploying robots. How should robots make decisions as they develop more autonomy? What should we do ethically and what ethical issues arise for us as we interact with robots? Should we give them any rights? We have a human rights code; will there be a robot rights code?

There are, then, three fundamental questions we have to address. First, what should we humans do ethically in designing, building, and deploying robots? Second, how should robots decide, as they develop autonomy and free will, what to do ethically? Third, what ethical issues arise for us as we interact with robots?

Asimov's Laws of Robotics. In considering these questions we will go back to Isaac Asimov (1950) as he was one of the earlier thinkers about these issues, who put forward some interesting, if perhaps naïve, proposals. His original three Laws of Robotics are: I. A robot may not harm a human being, or, through inaction, allow a human being to come to harm. II. A robot must obey the orders given to it by human beings except where such orders would conflict with the First Law. III. A robot must protect its own existence, as long as such protection does not conflict with the First or Second Laws.

Asimov's Answers. Asimov's answers to those questions I posed are that, first, you must put those laws into every robot, and by law manufacturers would have to do that. Second, robots should always have to follow the prioritized laws. He did not say much about the third question. His plots arise mainly from the conflict between what the humans intend the robot to do and what it actually does do, or between literal and sensible interpretations of the laws because they are not codified in any formal language. He discovered many hidden contradictions, but they are not of great interest here. What is of interest and important here is that, frankly, the laws and the assumptions behind them are naïve. That is not to blame Asimov, as he was very early and pioneered the area, but we can say that much of the ethical discussion nowadays remains naïve. It presupposes technical abilities that we just do not have yet.

What We Need

We do not currently have adequate methods for modeling robot structure and functionality, of predicting the consequences of robot commands and actions, and of imposing requirements on those actions, such as reaching the goal but doing it in a safe way and making sure that the robot is always live, with no deadlock or livelock. And most importantly, one can put those requirements on the robot, but one has to be able to find out whether the robot will be able to satisfy those requirements. We will never have, for real robots, 100 percent guarantees, but we do need within-epsilon guarantees. Any well-founded ethical discussion presupposes that we (and robots) do indeed have such methods. That is what we require.

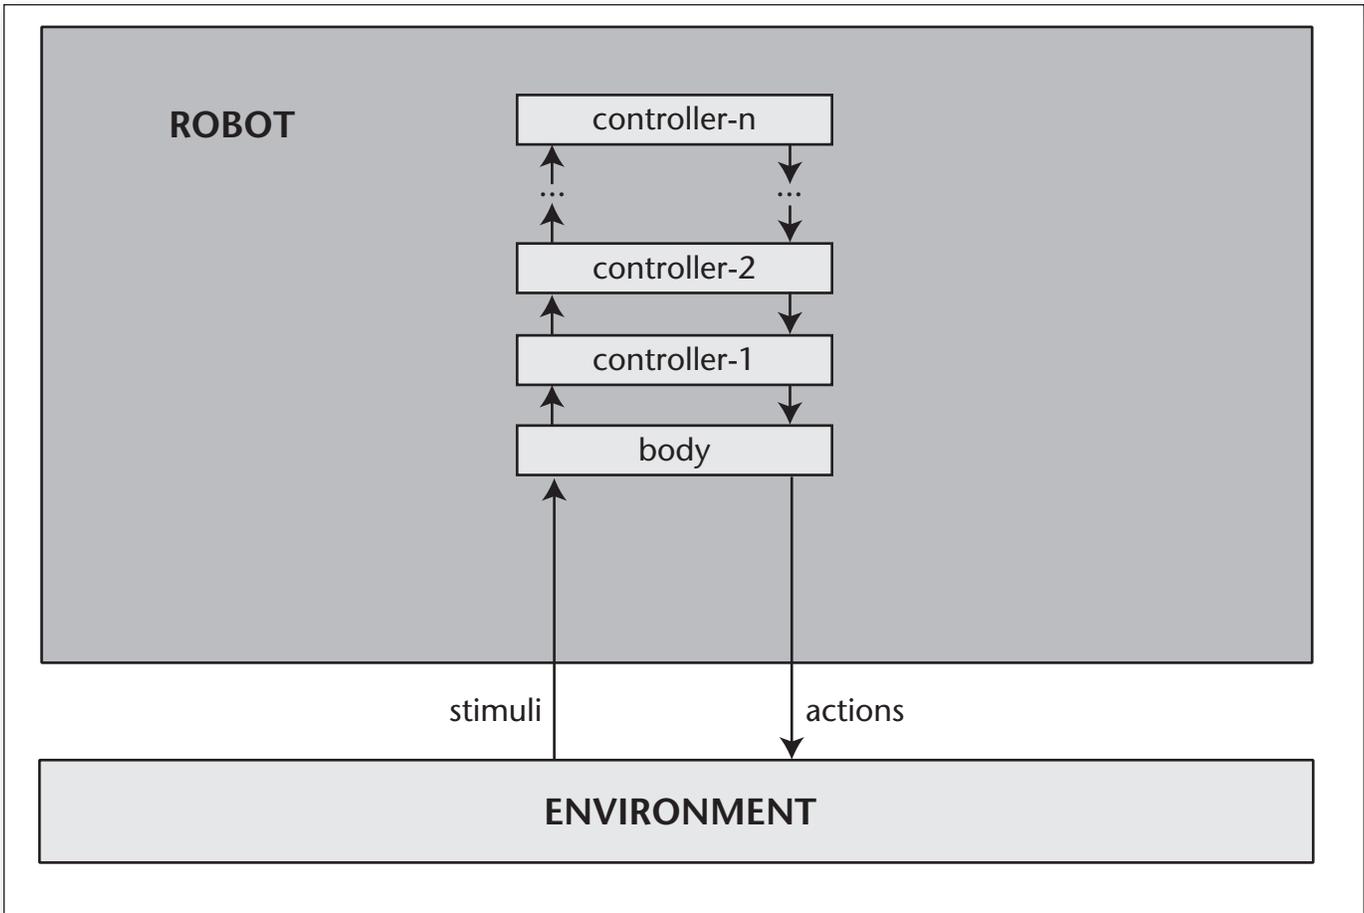


Figure 13. A Vertical Robotic System Architecture.

Theory Wanted

So, finally coming back to the constraint-based agent theory, it should help to satisfy those requirements. In short, we need a theory with a language to express robot structure and dynamics, language for constraint-based specifications, and a verification method to determine whether a robot described in the first language will be likely to satisfy its specifications described in the second language.

Robots as Situated Agents

What kind of robots, then, are we thinking about? These are *situated* robots tightly coupled to the environment; they are not universal robots. Remember *Rossum's Universal Robots* (Capek 1923)? We are not going to build universal robots. We are going to build—we are building—very situated robots that are functioning in particular environments for particular tasks. But those environments are, typically, highly dynamic. There are other agents. We have to consider social roles. There is a

very tight coupling of perception and action, perhaps at many different levels. We now know that the human perceptual system is not a monolithic black box that delivers a three-dimensional model from retinal images. There are many visual subsystems dealing with recognition, location, orientation, attention, and so forth. Our robots will be like that as well.

It is not “cheating” to embody environmental constraints by design, evolution, or learning. It was cheating in the old GOFAIR paradigm, which *did* aim at universal robots. Everything had to be described in, say, the logic, and one could not design environmental constraints into the robots. We think just following biology and natural evolution is the way to go, and learning will play a major part. Evolution is learning at the species level. Communication and perception are very situated. The architectures are online, and there is a hierarchy of time scales and time horizons. Critically, we want to be able to reason about the agent's correctness as well as in the agents. We do not require the agents to do reasoning—they may not—but

certainly we want to be able to reason about them. When we think back to the GOFAIR model, we never actually did that. The reasoning was in the agent's head alone, and we assumed that if it was correct everything else was correct. Finally, as I mentioned earlier, one cannot just graft a symbolic system on top of a signal-control-based system and expect the interface to be clean, robust, reliable, debuggable, and (probably) correct. So the slogan is “No hybrid models for hybrid systems.”

Vertical Architecture

To satisfy those requirements for situated agents we have to throw away the horizontal three boxes architectural model and move to a vertical “wedding cake” architecture. As shown in figure 13, as one goes up these controllers, each controller sees a virtual body below it, modularizing the system in that way. Each controller, as one goes higher, is dealing with longer time horizons but with coarser time granularity and different kinds of perception. Each controller will only know what it needs to know. This architectural approach was advocated by James Albus (1981) and Rodney Brooks (1986). It corresponds quite closely to biological systems at this level of abstraction.

A Constraint-Based Agent

We are interested in constraint-based agents. They are situated; they will be doing constraint satisfaction but in a more generalized sense, not in the GOFCS sense. These constraints may be prioritized. Now we conceive of the controller of the agent or robot as a *constraint solver*.

Dynamic Constraint Satisfaction

Consider the generalization of constraint satisfaction to *dynamic constraint satisfaction*. A soccer example will serve us.

Imagine a humanoid robot trying to kick a soccer ball. In figure 14, we can see the projection into a two-dimensional space of a complex phase space that describes the position and velocity of the limbs of the robot and the ball at time (t). Each flow line in the figure shows the dynamics of the evolution of the system from different initial conditions. The controller has to be able to predict where the robot should move its foot to, knowing what it knows about the leg actuators and the ball and where it is moving and how wet the field is and so on, to make contact with the ball to propel it in the right direction. That corresponds to the 45 degree line $y = x$. So x here is the ball position on the horizontal axis and y is the foot position on the vertical axis. That is the constraint we are trying to solve. If the controller ensures the dynamical system always goes to (or approaches, in the limit) that constraint and stays there, or maybe if it doesn't stay there but always returns to it soon enough,

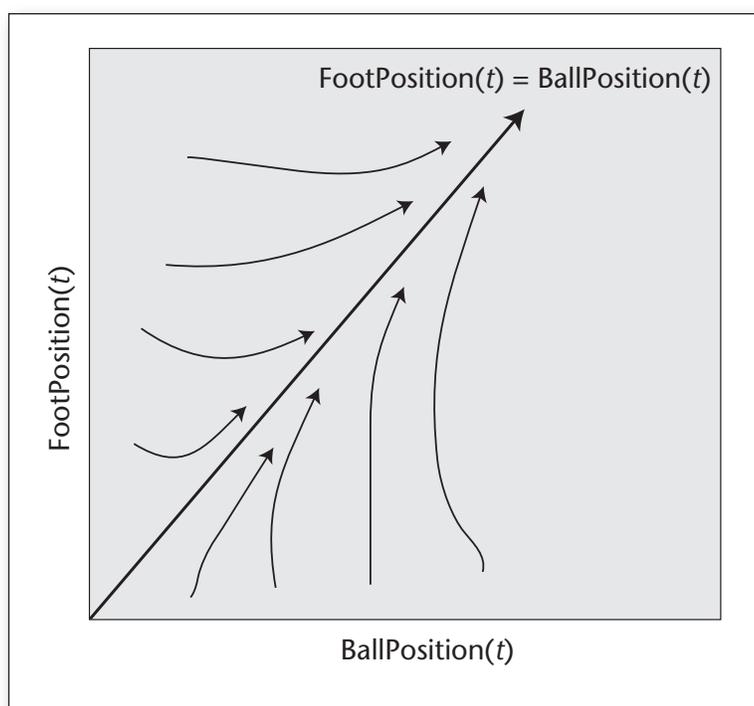


Figure 14. Solving a Dynamic Soccer Constraint.

then we say that this system is solving that constraint, $FootPosition(t) = BallPosition(t)$. In hybrid dynamic systems language, we say the coupled agent environment system *satisfies the constraint* if and only if the constraint solution set, in the phase space of that coupled hybrid dynamic system, is an *attractor* of the system as it evolves. Incidentally, that concept of online hybrid dynamic constraint satisfaction subsumes the entire old, discrete offline GOFCS paradigm (Zhang and Mackworth 1993).

Formal Methods for Constraint-Based Agents

The constraint-based agent (CBA) framework consists of three components: (1) Constraint Nets (CN) for system modeling; (2) timed for-all automata for behavior specification; and (3) model-checking and Liapunov methods for behavior verification.

These three components correspond to the tripartite requirement for the theory I said we wanted earlier. Ying Zhang and I developed these formal methods for constraint-based agents (Zhang and Mackworth 1995; Mackworth and Zhang 2003). First, there is an architecture for distributed asynchronous programming languages called Constraint Nets. Programs in CN represent the robot body, the controller, and the environment. In them are represented constraints that are local on the structure and dynamics of each system. For

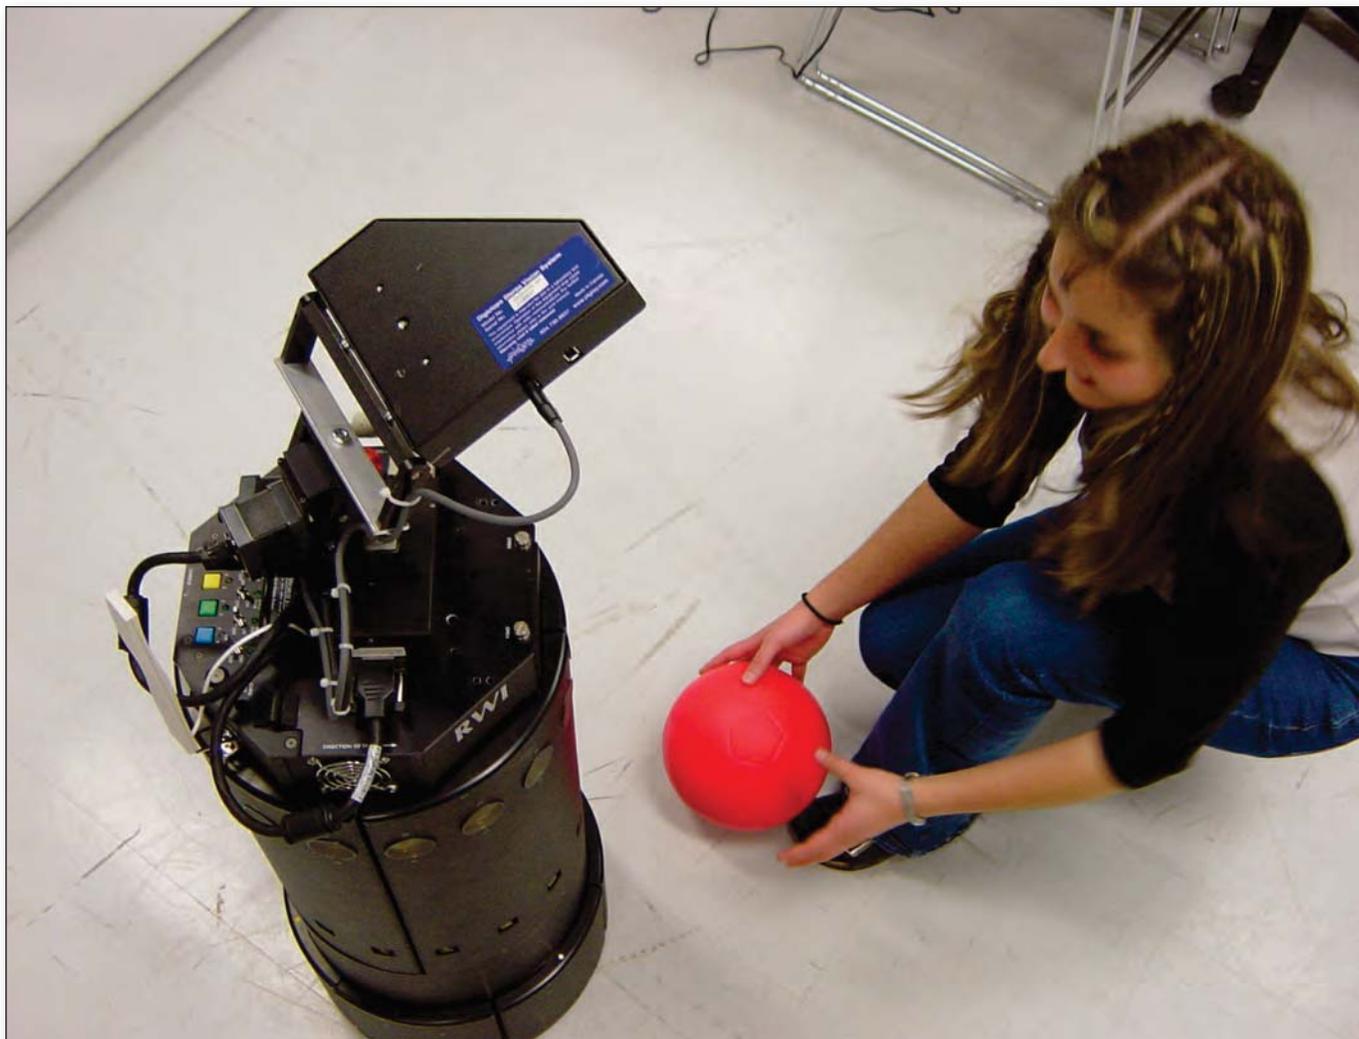


Figure 15. A Robot and a Human Play Ball.

behavior specification we use either temporal logics or timed for-all automata. For verification techniques we have used model checking or generalized Liapunov techniques taken from the standard control literature but generalized for symbolic as well as numerical techniques. Rather than present any technical detail here, I shall sketch a case study, again using soccer.

A Soccer Case Study with Prioritized Constraints

Suppose we want to build a robot soccer player that can move around the world, and repeatedly find, track, chase, and kick the soccer ball. The setup is shown in Figure 15. Pinar Muyan-Özçelik built a controller for this robot to carry out the task using the constraint-based agent methodology

(Muyan-Özçelik and Mackworth 2004). The detailed view of the robot in figure 16 shows a robot base, which can only move in the direction it is facing but can rotate in place to move in a new direction. There is a pan-tilt unit that serves as a neck and a trinocular color camera on top of it that can do stereo vision, but in this experiment we used monocular color images only.

This is a very simple, almost trivial example, but even here you get a rich complexity of interaction with emergent behavior. Imagine that you have got very simple controllers that can solve each of these constraints: (1) to get the ball in the image, (2) if the ball is in the image, center it, (3) make the base heading equal to the pan direction. Imagine that you are a robot and remember you can only move forward in the direction you are facing. If you turn your head to the left and acquire the ball in the image over there then you have to turn your

body to the left towards it, and as you are tracking the ball in the image, you have to turn your head to the right in the opposite direction. This is analogous to the well-known vestibuloocular reflex (VOR) in humans. Now you are looking at the ball and facing towards it, so now you can move towards it and hit the ball. The last constraint is for the robot to be at the ball. These are the constraints, and from that will emerge this unified behavior: acquire, track, chase, and kick the ball if we can satisfy these constraints, but in that priority order. If at any time one is satisfying a lower priority constraint and a higher priority constraint becomes unsatisfied, the controller must revert to resatisfying it.

The prioritized constraints are: Ball-In-Image (I), Ball-In-Center (C), Base-Heading-Pan (H), Robot-At-Ball (A). The priority ordering is: $I > C > H > A$. We want to satisfy those prioritized constraints. That is the specification. The specification for this system is that one has to solve those four constraints with that priority. That is all one would say to the system. It is a declarative representation of the behavior we want the system to exhibit. We can automatically compile that specification into a controller that will in fact exhibit that emergent behavior. We can conceptualize these prioritized constraint specifications as generalizations of the GOFAIR linear sequence plans.

Constraint-Based Agents in Constraint Nets

Suppose we are given a prioritized constraint specification for a controller at a certain level in the controller hierarchy as in figure 17. The specification involves Constraint-1, Constraint-2, and Constraint-3. It requires this priority order: Constraint-1 > Constraint-2 > Constraint-3.

We assume we have a simple solver for each constraint, Constraint Solver-1, Constraint Solver-2, and Constraint Solver-3. Constraint-1 is the highest priority, so if it is active and not satisfied, its solver indicates, "I'm not satisfied now; I'd like you to do this to satisfy Constraint-1." It might be a gradient descent solver, say. Its signal would go through Arbiter-1. The arbiter knows this is higher priority and its signal passes it all the way through to Arbiter-2 as well, to the motor outputs. If Constraint-1 is satisfied, Arbiter-1 will let Constraint Solver-2 pass its outputs through and so on. If there is any conflict for the motors, that is how it is resolved. If there is no conflict then the constraints can be solved independently because they are operating in orthogonal spaces. Using that architecture we built a controller for those constraints in the soccer player, and we tested it all in simulation, and it works. It works in a wide variety of simulated conditions; it works with the same controller in a wide variety of real testing situa-

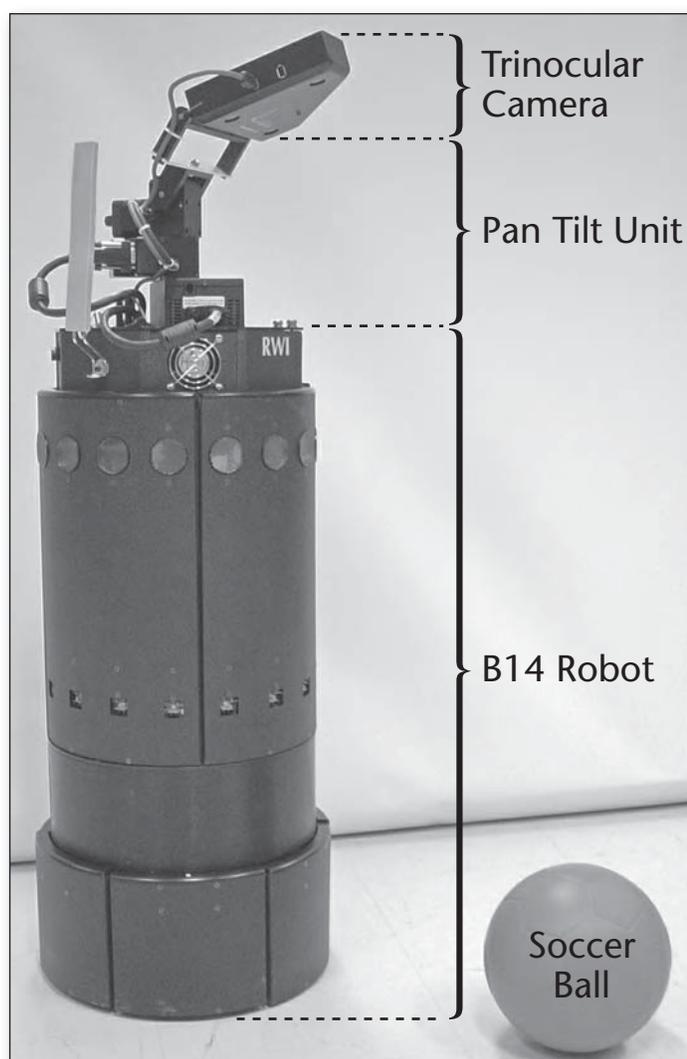


Figure 16. A Simple Soccer Player.

tions. We found that the robot always eventually kicks the ball repeatedly, both in simulation and experimentally. In certain circumstances, we can *prove* that the robot always eventually kicks the ball repeatedly. We conclude that the constraint-based agent approach with prioritized constraints is an effective framework for robot controller construction for a simple task.

Just as an aside, this is a useful way to think about a classic psychology problem. Carl Lashley, in a seminal paper called "The Problem of the Serial Ordering of Behavior" (Lashley 1951), was writing about speech production, but the sequencing problem arises for all behaviors. Suppose one has a large number of goals one is attempting to satisfy. Each goal is clamoring to be satisfied. How do the goals get control of the actuators? How does one sequence that access? How does one make sure it is

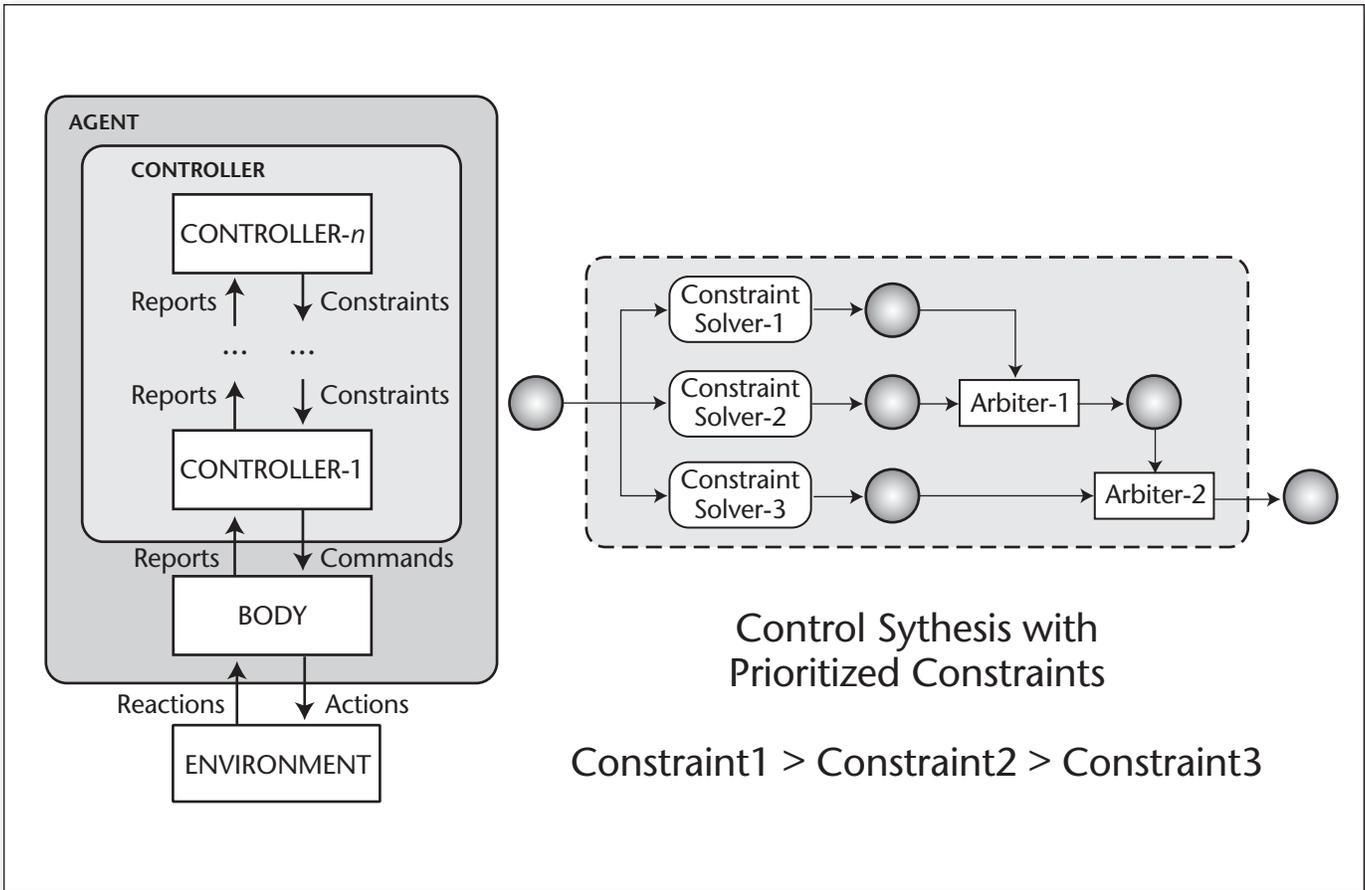


Figure 17. Synthesizing a Controller from a Prioritized Constraint Specification.

robust? For prioritized constraints, it is robust in the sense that if the controller senses the loss of satisfaction of a higher priority constraint it will immediately resume work on resatisfying it. One can also think of what we have done as a formalization of subsumption (Brooks 1991). So this is a solution; I am not saying it is *the* solution. It is a very simple-minded solution, but it is a solution to the classic problem of the serial ordering of behavior. And it demonstrates that prioritized constraints can be used to build more reliable dynamic agents.

Modeling Uncertainty

So far, nothing has been said about the element of chance but, of course, in real robots with real environments there will be much noise and uncertainty. Robert St-Aubin and I have developed probabilistic constraint nets using probabilistic verification (St-Aubin, Friedman, and Mackworth 2006). As shown in Figure 18, there will be uncertainty in the model of the dynamics and in the dynamics themselves and in the robot’s measure-

ment of the world. Moreover, one will be unable to fully model the environment, so there will be environmental disturbances.

Observations

Stepping back we observe that a very simple idea, constraint satisfaction, allows us to achieve intelligence through the integration of proactive and responsive behaviors and it is uniform top to bottom. We see in the formal prioritized framework the emergence of robust goal-seeking behavior. I propose it as a contribution to the solution of the problem of a lack of technical foundation to many of the proposals for robot ethics. So if one asks, “Can robots do the right thing?” The answer so far is, “Yes, sometimes they can do the right thing, almost always, and we can prove it, sometimes.”

Challenges for AAI

As noted throughout, evolution is a dynamic process. I have discussed how constraint satisfaction has changed, how agents change, how our

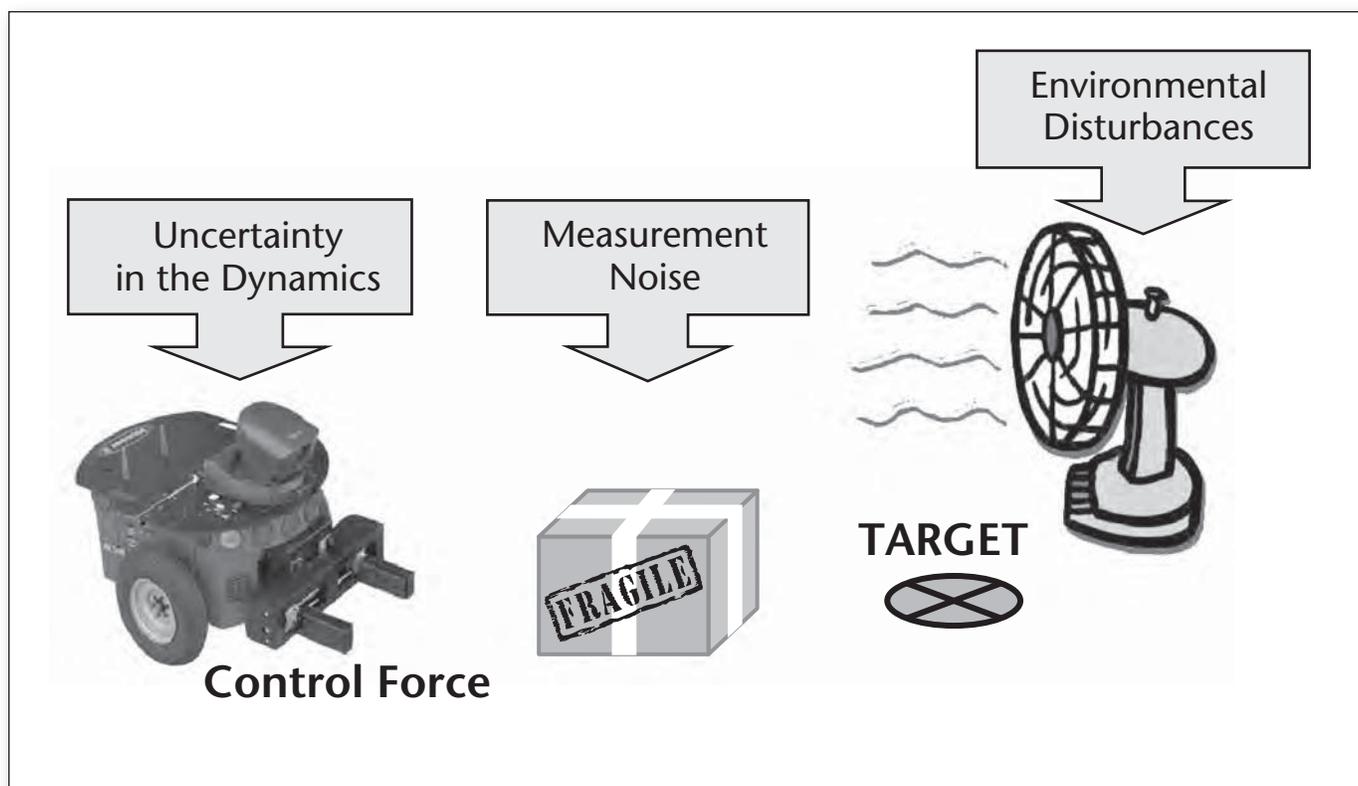


Figure 18. Uncertainty in Robots.

model of agents has changed, and how AI is changing as we dropped the metaassumptions that held us back in GOFAIR and evolved to the situated agent model.

AAAI, as an organization, is an agent; it is operating in an environment, a global environment. It is a very competitive but also collaborative environment. In fact, to me, it is analogous to soccer. Of course, apparently everything looks like soccer to me, so that is not surprising. Think of AAAI as a soccer team. Keep your eye on the ball, know what game is going on, and know what the competition is doing. We constitute the premier AI membership society on the planet. I say this with pride. It is healthy; we have great involvement from our community. We see at our conferences the excitement generated by all the different events, competitions, and papers, and participant interaction.

AAAI is actually changing AI, and that is what we should be doing through all our activities including the workshops and symposia. And we are. But the world is also changing; in fact, AI has changed in a way that maybe we have not recognized as a society. AI does now have a level of scientific and engineering maturity. In some of our early days we rather pushed aside control theory and pushed aside pattern recognition and the like.

We insisted we were doing “important symbolic work” and dismissed other aspects of work that we now value. We have to change that attitude, and we have to change it in our society and in our field. We are now doing so more and more. Now we rejoice in a new AI. We recognize that uncertainty is critically important thanks to Peter Cheeseman (1985), Judea Pearl (1988), and many others, and we are including fields such as statistics, machine learning, pattern recognition, operations research, control, neuroscience, and psychology and welcome them all. The beauty of the way we work as a field is that we are problem focused not discipline focused, so we do not have those disciplinary silo traps to worry about. AI is now everywhere.

Globalization has hit us too. We came to realize that although we called ourselves the American Association for Artificial Intelligence, about one-third of our members are not American. About half of the conference papers coming in and the attendees at our conferences are not American. Do they feel welcome? Of course, for a Canadian president to ask that question was a little tricky, but I was assured by my American colleagues that it was a critical question to ask. Thus, we decided to “go for it” to recognize the current reality. In proposing to change the name of the society, we were anticipat-

ing a bit of a firestorm, which, happily, did not happen. Certainly our membership, over 90 percent, was enthusiastically in favor of it, and, interestingly, many of the international societies were in favor of changing our name to the Association for the Advancement of Artificial Intelligence, which recognizes the current nature of our global situation.

We still must remember we have three roles: national, regional, and international. We have a national role; we have to represent our interests in scientific advancement in Washington and to be part of the Computing Research Association looking at the science and engineering funding situation in the United States and in Canada; and we also have both a regional role and, increasingly, an international role. We need to communicate better with our sibling AI and computer-science societies and push harder for global coordination.

Strategic Directions for AAAI

Given the rapid change in AI and in the environment for scientific research and global communications, this is an opportune time to reflect upon strategic directions for AAAI. AAAI, until 2007, had never had a formal strategic plan. Thus, the Strategic Planning Board and Council started a planning process with a Strategic Planning Working Group in 2007, which is now well underway. We need to consider our vision, our mission, our values, and our goals as a scientific society. These goals include supporting and facilitating the basic and applied science of AI, enabling and facilitating the practice of AI, conducting outreach and advocacy, and being an agile and responsive professional society. Some of the key issues that have arisen include education for various audiences; communication forums (conferences, workshops, symposia); open access for publications, knowledge, and data; ensuring we provide tangible member benefits; building community; exploiting media; and building collaborative linkages with other AI societies, conferences, and journals.

We have done some good work in education with AI Topics, led by Bruce Buchanan and Jon Glick, but we could do more to encourage AI curriculum development. Clearly, AI is the most exciting field in computer science, but most undergraduate programs do not offer it until third year. Perhaps we could help with enrollment and recruitment if we could make it more accessible earlier in the curriculum. Recent AAAI meetings on AI and education are promoting and enabling this work.

Open access to publications is a problem and opportunity we do have to address. The publications access task force under president-elect Martha Pollack is carefully considering the options, acknowledging that access to publications is a significant benefit to our members. The *AI Mag-*

azine, guided by David Leake, does a wonderful job. New forms of publication, such as video on the web, are now being actively encouraged. We are also archiving and preserving historical AI videos.

Finally, AAAI has a leadership role to play in the international AI community. In working cooperatively with all our sibling AI societies and conferences, we need more coordination and coherence, internationally. Obviously, this must be done carefully, in a collegial fashion, but it must be done.

How Can We Help?

As I contemplated what I wished to address in concluding, I realized I wanted to touch upon a few of my more personal thoughts. AI is now mature, both as a science and in its technologies and applications as an engineering discipline. There are themes of AI in service to people and to our planet; themes that we should explore. We no longer have the luxury of ignoring the global crises of global warming, poverty, food production, sustainability, arms control, health, education, the aging population, and demographic issues. We do have a unique perspective and the skills to contribute practically to addressing many of these concerns. There is an opportunity to join with people who are already dealing with these issues, not in an arrogant way but collaboratively. We could, as an example, improve access to tools for learning about AI so that people could be empowered to try our techniques on their problems. We know we cannot go in as experts telling people how to solve their problems, but we can make tools available so that someone already knowledgeable in a domain might realize that machine learning, say, offers radically new and potentially better approaches to a problem. The success of RoboCup and other games and competitions shows how effective they can be as learning and teaching environments. The One Laptop Per Child (OLPC) initiative has already made enormous strides; we could be supporting this more actively. Clearly, there is ample room for imaginative thinking in this area.

There are many opportunities for AI to have a positive impact upon the planet's environment. It is tremendously encouraging that work for betterment has already been done or is being contemplated. We have already considered some of the environmental impact of intelligent cars and smart traffic control. All the work on combinatorial auctions, already applied to spectrum allocation and logistics, could be further applied to optimizing energy supply and demand and to carbon offsets. There could be much more work on smart energy controllers with distributed sensors and actuators that would improve energy use in buildings. We could perhaps use qualitative modeling techniques for climate scenario modeling.

Finally, assistive technology for the disabled and aging is being pioneered by researchers such as Martha Pollack (2005) and Henry Kautz (Liu et al. 2006). We need to listen very carefully to what they are telling us. There is a huge opportunity for us, an opportunity that is extremely attractive to graduate students. One issue, which has been mentioned earlier, is the progress-hindering question of liability, which can be, and must be, addressed. Assisted cognition is one application, but work is also proceeding in assisted perception and assisted action, for example, in the form of smart wheelchairs and companions for the elderly and nurses' assistants in long-term care facilities, following the lead of Alex Mihailidis and Jesse Hoey (Mihailidis et al. 2007). All of the research topics mentioned will, of course, have many obstacles to overcome. It is an exciting time to do research in AI and its applications. Perhaps some of these will serve as challenges to our younger, and not so young, researchers.

Conclusion

To recap, I have sketched the evolution of constraint satisfaction, of agents, of models of agents, of AI, and of AAI. Given the scope of those topics, this sketch has necessarily been idiosyncratic, personal, and incomplete. However, I hope to have persuaded you of the thesis that constraint satisfaction is central to intelligent behavior.

Acknowledgements

I am most grateful to all of the students, colleagues, and collaborators who have contributed to some of the work mentioned here: Saleema Amershi, Rod Barman, Le Chang, Cristina Conati, Giuseppe Carenini, Pooyan Fazli, Gene Freuder, Joel Friedman, Stewart Kingdon, Jaçek Kisynski, Byron Knoll, Jim Little, David Lowe, Valerie McRae, Jefferson Montgomery, Pinar Muyan-Özçelik, Dinesh Pai, David Poole, Michael Sahota, Fengguang Song, Robert St-Aubin, Pooja Viswanathan, Bob Woodham, Suling Yang, Ying Zhang, and Yu Zhang. The support of Carol Hamilton and all her staff made it easier to serve as AAI president. David Leake helped with this article. I am also grateful to those who served on the AAI Council and on the Executive Committee and all the other volunteers who make AAI special. Funding was provided by the Natural Sciences and Engineering Research Council of Canada and through the support of the Canada Research Chair in Artificial Intelligence.

Notes

1. Alspace (formerly known as Cspace) is an environment for learning and teaching AI featuring animations of many AI algorithms (Amershi et al. 2008, Knoll et al. 2008). The Alspace site is at www.Alspace.org.
2. William Shakespeare, *Hamlet*, Act 1.

3. www.cs.ubc.ca/~mack/RobotSoccer.htm.
4. www.cs.ubc.ca/~mack/RobotSoccer.htm.
5. www.fira.net/soccer/mirosot/overview.html.
6. www.tartanracing.org, cs.stanford.edu/group/roadrunner.

References

- Albus, J. S. 1981. *Brains, Behavior and Robotics*. New York: McGraw-Hill.
- Amershi, S.; Carenini, G.; Conati, C.; Mackworth, A. K.; and Poole, D. L. 2008. Pedagogy and Usability in Interactive Algorithm Visualizations: Designing and Evaluating Cspace. *Interacting with Computers*, 20(1): 64–96.
- Anderson, M., and Anderson, S. L. 2007. Machine Ethics: Creating an Ethical Intelligent Agent. *AI Magazine*, 28(4): 15–26.
- Asimov, I. 1950. *I, Robot*. New York: Gnome Press.
- Brooks, R. A. 1986. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation* RA-2(1): 14–23.
- Brooks, R. A. 1991. Intelligence without Reason. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, 569–595. San Mateo, CA: Morgan Kaufmann.
- Capek, K. 1923. *R.U.R. (Rossum's Universal Robots): A Fantastic Melodrama in Three Acts and an Epilogue*. Garden City, NY: Doubleday.
- Cheeseman, P. 1985. In Defense of Probability. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 1002–1009. San Mateo, CA: Morgan Kaufmann.
- Dresner, K., and Stone, P. 2008. A Multiagent Approach to Autonomous Intersection Management. *Journal of Artificial Intelligence Research* 31: 591–656.
- Feuerbach, L. A. 1854. *The Essence of Christianity*. London: John Chapman.
- Freuder, E. C., and Mackworth, A. K. 2006. Constraint Satisfaction: An Emerging Paradigm. In *Handbook of Constraint Programming*, ed. F. Rossi, P. van Beek, and T. Walsh, 13–28. Amsterdam: Elsevier.
- Haugeland, J. 1985. *Artificial Intelligence: The Very Idea*. Cambridge, MA: The MIT Press.
- Hsu, F. 2002. *Behind Deep Blue: Building the Computer That Defeated the World Chess Champion*. Princeton, NJ: Princeton University Press.
- Kitano, H., ed. 1998. *RoboCup-97: Robot Soccer World Cup I*. Lecture Notes in Computer Science 1395. Heidelberg: Springer.
- Knoll, B.; Kisynski, J.; Carenini, G.; Conati, C.; Mackworth, A. K.; Poole, D. L. 2008. Alspace: Interactive Tools for Learning Artificial Intelligence. In *Proceedings of the AAI AI Education Colloquium*. Technical Report WS-08-02, Menlo Park, CA: AAI Press.
- Lashley, K. S. 1951. The Problem of Serial Order in Behavior. In *Cerebral Mechanisms in Behavior*, ed. L. A. Jeffress, 112–136. New York: Wiley.
- Lennon, J. 1980. Beautiful Boy (Darling Boy). Song lyrics. On album *Double Fantasy*. New York: Geffen Records.
- Liu, A. L.; Hile, H.; Kautz, H.; Borriello, G.; Brown, P. A.; Harniss, M.; and Johnson, K. 2006. Indoor Wayfinding: Developing a Functional Interface for Individuals with

The Eighth and Ninth International
Conferences on Autonomous
Agents and Multiagent Systems

aamas 2009

Budapest, Hungary, 10 - 15 May 2009

aamas 2010

Toronto, Canada 10 - 14 May 2010

AAMAS is the leading scientific conference for research in autonomous agents and multiagent systems.

The conference provides a single, high-profile, internationally respected forum for research in the theory and practice of autonomous agents and multiagent systems, and includes a wide range of advanced tutorials and workshops as part of its programme.

www.aamas-conference.org

Cognitive Impairments. In *Proceedings of the 8th International ACM SIGACCESS Conference on Computers and Accessibility*, 95–102. New York: Association for Computing Machinery.

McLuhan, M. 1964. *Understanding Media: The Extensions of Man*. New York: New American Library.

Mackworth, A. K. 1977. Consistency in Networks of Relations. *Artificial Intelligence* 8(1): 99–118.

Mackworth, A. K. 1993. On Seeing Robots. In *Computer Vision: Systems, Theory and Applications*, eds. A. Basu and X. Li, 1–13. Singapore: World Scientific Press.

Mackworth, A. K., and Zhang, Y. 2003. A Formal Approach to Agent Design: An Overview of Constraint-Based Agents. *Constraints* 8(3): 229–242.

Maxwell, J. C. 1868. On Governors. In *Proceedings of the Royal Society of London*, 16, 270–283. London: Royal Society.

Mihailidis, A.; Boger, J.; Candido, M.; and Hoey, J. 2007. The Use of an Intelligent Prompting System for People with Dementia. *ACM Interactions* 14(4): 34–37.

Miller, G. A.; Galantner, E.; and Pribram, K. H. 1960. *Plans and the Structure of Behavior*. New York: Holt, Rinehart, and Winston.

Mori, M. 1982. *The Buddha in the Robot*. Tokyo: Charles E. Tuttle Co.

Muyan-Özçelik, P., and Mackworth, A. K. 2004. Situated Robot Design with Prioritized Constraints. In *Proceedings of the International Conference on Intelligent Robots and Sys-*

tems (IROS 2004), 1807–1814. Los Alamitos, CA: IEEE Computer Society.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA: Morgan Kaufmann.

Pollack, M. E. 2005. Intelligent Technology for an Aging Population: The Use of AI to Assist Elders with Cognitive Impairment. *AI Magazine* 26(2): 9–24.

Reiter, R. 1978. On Closed World Data Bases. In *Logic and Data Bases*, eds. H. Gallaire and J. Minker, 119–140. New York: Plenum.

Rossi, F.; van Beek, P.; and Walsh, T., eds. 2006. *Handbook of Constraint Programming*. Amsterdam: Elsevier Science.

Sahota, M., and Mackworth, A. K. 1994. Can Situated Robots Play Soccer? In *Proceedings of Artificial Intelligence '94*, 249–254. Toronto, ON: Canadian Society for Computational Studies of Intelligence.

Shelley, M. W. 1818. *Frankenstein; or, The Modern Prometheus*. London: Lackington, Hughes, Harding, Mavor, and Jones.

Somerville, M. 2006. *The Ethical Imagination: Journeys of the Human Spirit*. Toronto: House of Anansi Press.

St-Aubin, R.; Friedman, J.; and Mackworth, A. K. 2006. A Formal Mathematical Framework for Modeling Probabilistic Hybrid Systems. *Annals of Mathematics and Artificial Intelligence* 37(3–4): 397–425.

Thrun, S. 2006. Winning the DARPA Grand Challenge. Invited Talk presented at the Innovative Applications of Artificial Intelligence conference (IAAI-06), Boston, Massachusetts, July 16–20.

Visser, U., and Burkhard, H.-D. 2007. RoboCup: 10 Years of Achievements and Challenges. *AI Magazine* 28(2): 115–130.

Waltz, D. L. 1975. Understanding Line Drawings of Scenes with Shadows. In *The Psychology of Computer Vision*, ed. P. H. Winston, 19–92. New York: McGraw-Hill.

Zhang, Y., and Mackworth, A. K. 1993. Constraint Programming in Constraint Nets. In *Proceedings of the First Workshop on Principles and Practice of Constraint Programming*, 303–312. Padua: Association for Constraint Programming.

Zhang, Y., and Mackworth, A. K. 1995. Constraint Nets: A Semantic Model for Dynamic Systems. *Theoretical Computer Science* 138: 211–239.

Alan K. Mackworth is a professor of computer science, Canada Research Chair in Artificial Intelligence, and the founding director of the Laboratory for Computational Intelligence at the University of British Columbia. He received a B.A.Sc. (Toronto, 1966), A.M. (Harvard, 1967), and D.Phil. (Sussex, 1974). He has worked primarily on constraint-based artificial intelligence with applications in vision, constraint satisfaction, robot soccer, and constraint-based agents. He has served as chair of the IJCAI Board of Trustees and as president of the Canadian Society for Computational Studies of Intelligence and the Association for the Advancement of Artificial Intelligence (AAAI); he is now past-president of AAAI.