

**Toward a More Intelligent and Cooperative
Generation of Construction Robots**

by

Boyd C. Paulson

TECHNICAL REPORT

Number 22B

January, 1992

Stanford University

Copyright © 1992 by
Center for Integrated Facility Engineering

If you would like to contact the authors please write to:

*c/o CIFE, Civil Engineering,
Stanford University,
Terman Engineering Center
Mail Code: 4020
Stanford, CA 95305-4020*

Toward a More Intelligent and Cooperative Generation of Construction Robots¹

Boyd C. Paulson
Charles H. Leavell Professor of Civil Engineering
Department of Civil Engineering
Stanford University

Stanford, California 94305-4020
USA

ABSTRACT

In future construction environments, intelligent machines must harness considerable knowledge to plan and control their tasks in spite of being limited in their own pre-defined knowledge. This paper first briefly describes the achievements and limitations of today's construction robots, then addresses the more complex environment in which their successors will work. To better integrate robots coming from various sources, they will need a unifying core of intelligent software and a framework to communicate knowledge about designs and field operations. Tomorrow's machines will need better abilities to access knowledge sources in their environment and to work cooperatively with other agents. While the development of such capabilities is among the most challenging tasks facing robotics researchers, it is not too soon to start coming to a consensus about the general nature of the intelligence robots will need to integrate into their environment and become more productive under complex conditions.

¹Note. This paper first appeared as: Paulson, Boyd C., "Toward a More Intelligent and Cooperative Generation of Construction Robots," Keynote Paper, *Proceedings of the First Symposium on Construction Robotics*, Tokyo, Japan, June 25-27, 1990, pp. 1-10.

INTRODUCTION

There can be little doubt that, while construction was once considered a rather backward industry, it is increasingly being recognized for providing some of the most complex and interesting challenges to researchers working in computer science, robotics, artificial intelligence, sensors, communications and equipment engineering. For example, last March on the Stanford University campus, the annual meeting of the American Association for Artificial Intelligence (AAAI) was held almost simultaneously with the annual symposium of the Center for Integrated Facility Engineering (CIFE). The latter is a new research center, with substantial support from the Japanese construction industry among others, that is focused on the needs for better integration across the engineering design and construction process. Although the AAAI and CIFE meetings were not organized as a joint conference—indeed, they met in separate locations on campus—it was interesting to see that many of the AI and robotics researchers, including pioneers of AI and expert systems such as Edward Feigenbaum,² spent much of their time at the CIFE symposium, and some of the leading construction researchers such as Kiyoshi Niwa¹¹ of Hitachi attended the AAAI meeting. There was a genuine intellectual exchange among these people.

Although construction technology may indeed have evolved slowly for many decades until the early 1980's, the last few years mark the beginning of a renaissance in this field, and there is reason for optimism about the future. The pioneering work on construction robotics that has been emanating mostly from Japan is one of the best examples of this renaissance,⁵ and it is thus all the more of a privilege to be invited to address this audience.

Even as today's first generation of construction robots begins moving from the research laboratories out onto field construction projects, it already is time to be thinking about the future. Most of the successes thus far have been with very specialized machines working almost in isolation from other machines and workers around them, but this is not typical of most of the work that goes on in construction. My main focus in this paper will be on the need to better integrate future construction robots into the total construction environment, and with the engineering, design and administrative functions that support construction.

TODAY'S ROBOTS: ACCOMPLISHMENTS AND LIMITATIONS

Given the difficult and complex environment of construction, it is remarkable that robots are already performing routine tasks on some jobsites. The first construction robots have either been derived by adding sensors and computer-based controls to existing construction equipment (e.g., to control the cutting edges or screeds on various types of earthmoving and paving equipment), by adapting the comparatively rigid factory-type robots to

construction (e.g., for spraying fireproofing material or painting), or by developing hybrids of the two (e.g., robot arms mounted on tunnel machines). While the sophistication of their mechanisms and sensors has often been quite high, these robots have had only the most rudimentary forms of programmed "intelligence," and some machines that have been called robots are really just teleoperated devices without any programmed automation at all.

Most of the construction robots developed to date are stand-alone devices designed to perform narrowly defined tasks without the need to communicate or cooperate with other machines. The concept of a construction "crew" does not really apply yet to construction robots. However, coordinated teams of robots quite commonly perform sequential operations on factory assembly lines, and there are some formal communication mechanisms linking them together; it is only a matter of time before similar technology also moves to construction.

The challenges of the construction jobsite are much greater, however, than those of most factories.¹² To begin, the products of construction are much more complex and vastly larger. Furthermore, in contrast to the repetitive products that flow down production lines, the design of the construction product and the process to build it are usually uniquely adapted in each case. While the manufacturing process is largely steady-state once the production stage starts, that in construction is ever changing. The physical environment of construction is often much more hostile to machines as well as people, so machine design must account for extremes of weather, dust and unexpected forces. While these differences may seem obvious to researchers coming from a construction background, it is sometimes surprising that their implications are not at first obvious to researchers coming from other fields. The increasing cooperation among researchers from many disciplines is improving understanding at all levels, and more and more useful results can be expected.

While there are many challenges facing the advancement of construction automation and the development of more capable construction robots, perhaps the most difficult is that of developing the intelligent software to integrate future machines into the complex environment where they will work. Let us first take a closer look at the environment of construction from the intellectual perspective of a future construction robot.

THE FUTURE ENVIRONMENT FOR CONSTRUCTION ROBOTICS

The knowledge environment in which future construction robots are to be embedded is potentially vast. The environment is relatively unstructured and fast-changing, so there may not be accurate correspondence of an agent's knowledge about the environment to its real state at any time. In future construction field environments, intelligent machines, like their human counterparts, will thus need to harness considerable knowledge to plan and

control autonomous tasks in spite of the fact that they will be limited in their own knowledge and abilities. They will need a unifying core of intelligent software and a framework for defining and communicating knowledge about designs and field operations in a way that can effectively be utilized for their production tasks. At Stanford and elsewhere, researchers are working toward such a core of software in order to support the work of others on practical applications of robots in field conditions.^{7,13,14} Some of this research focuses on cognitive aspects of future machines to endow them with common modes of "thinking" and communicating—in effect a common *culture*—so that they can work together and with humans in groups.

The scope of research needed to build theories and core software to support construction robotics is also vast. Each step in this research should lead toward a general architecture handling the knowledge an agent needs to function productively in a knowledge environment. The resulting software could then be extended by developers of applications-oriented robots to handle particular areas of expertise, whether in managing other machines or in doing specific physical tasks. The ultimate objective should be to design and develop the general theory and software core for machine agents—the rudimentary "brains" of the beasts—which can then be embedded in agents specialized for particular tasks.

Figure 1 shows a broad conceptual view of the construction knowledge environment, and some ways in which the core software of an intelligent robot might interact with the environment. It illustrates the organizational context in which the robots might be working, the interfaces to computer-aided design (CAD) databases and reasoning, interactions with other field agents—both human and machine—and interfaces to knowledge sources in the world beyond the field. To establish such an environment with robotic *hardware* is not feasible now. Even if one could afford them, robots with sufficient flexibility and computing power for diverse operations do not yet exist. It is not too early, however, to begin coming to some form of consensus—perhaps even to be thinking about standards—so that such core software can evolve along with the sensors, control systems and mechanisms for these machines. I will outline some general needs and directions in the next section.

COMPONENTS OF ROBOTIC INTELLIGENCE

Before considering what should go into the core of construction robot software, it is important to think about some bounds on this software. For the most part we are looking at the intelligence to support the successful execution of construction tasks, and this probably will not require much in the way of, for example, aesthetic appreciation, the ability to sing or dance, or innumerable other facets that characterize human existence. Relative to the intelligence and human dimensions of a typical construction worker, we are

still looking at a most rudimentary kind of "intelligence" to form the core of a construction robot's software.

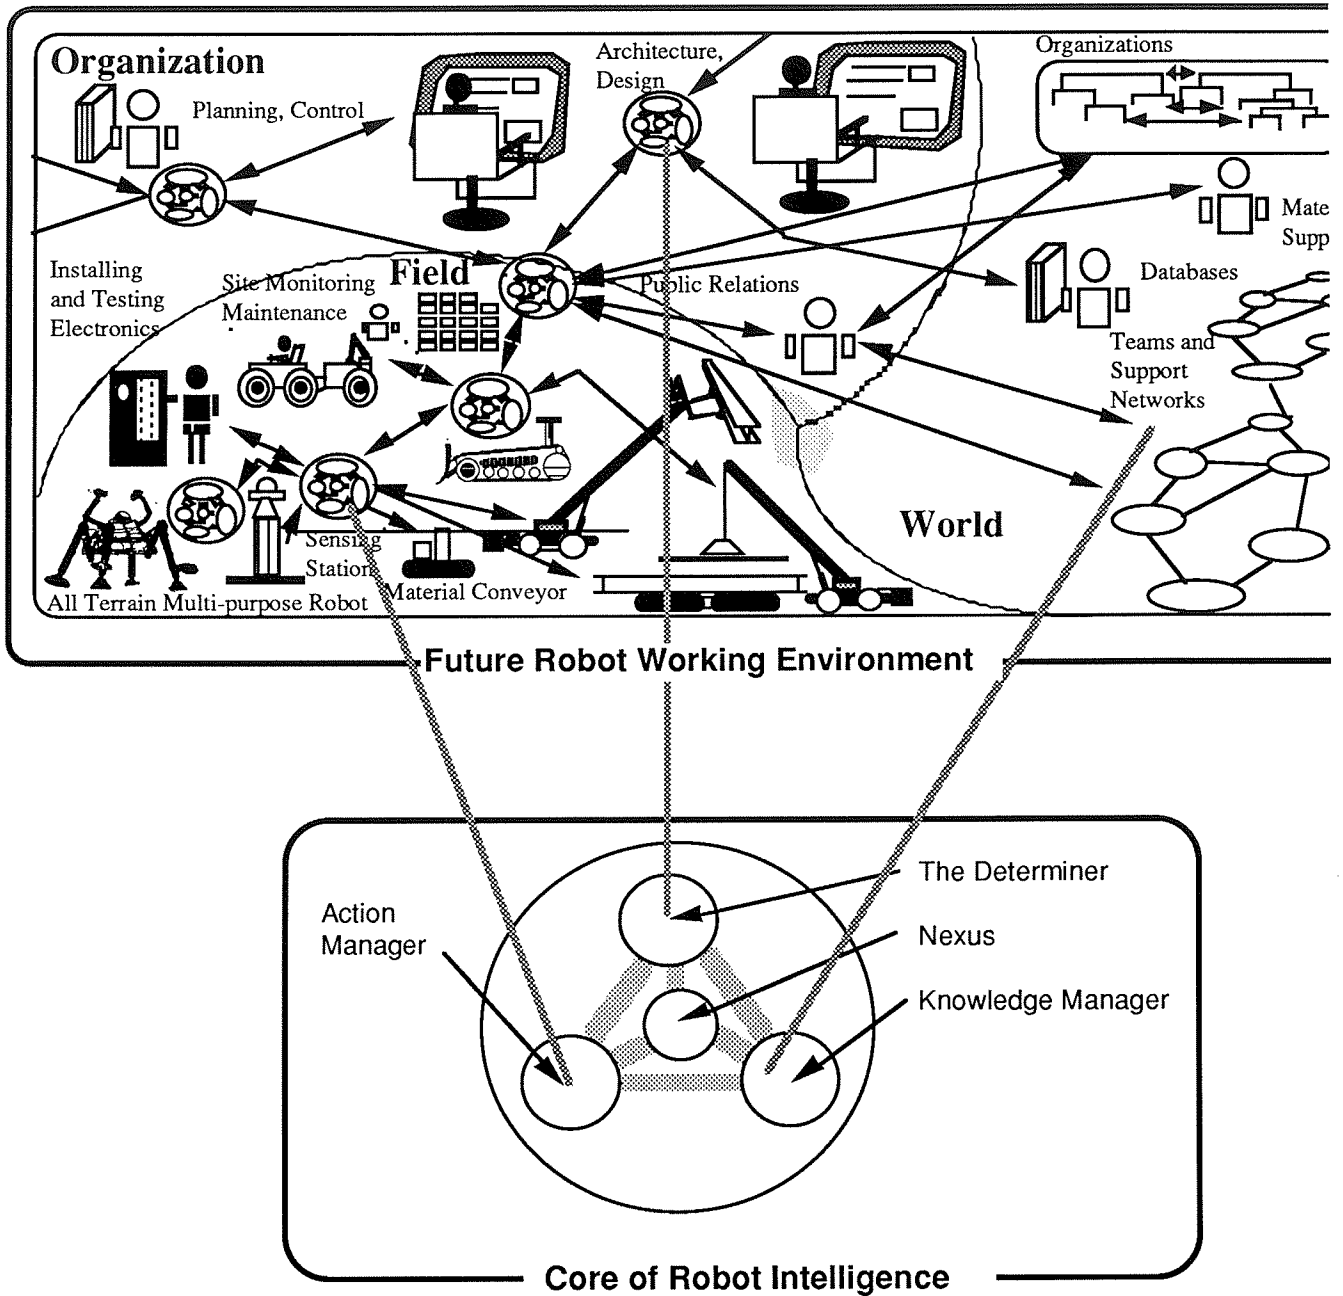


Figure 1—Conceptualization of Agent Core Software Interactions with the Knowledge Environment for Construction Automation

In general, what is needed is some way of modelling *within* robot agents some “understanding” of their environment, such as key characteristics of objects and other agents, in ways useful for reasoning. Among other things, researchers should seek to reduce the knowledge that needs to be encoded in machine systems *a priori* by enabling them to tap the vast knowledge sources in their environment when needed. This is extensibility, which some might call a simple form of learning. Automata should be able to assemble knowledge and enlist other agents needed to perform a task and respond dynamically to change. For example, robot reasoning and control software should deal with unexpected obstacles, road conditions, failure of a machine positioning system, damaged material, improper tools, or imprecise instructions.

Examples of some basic types of “cultural knowledge” that would be common to and not task-specific for such robots could include the following:

- Knowledge and abilities to deal with space and time, such as:
 - interpretation of and reasoning about geometric 3-D space
 - motion planning
 - Newtonian mechanics (velocity, forces, time, etc.)
 - the ability monitor the location and status of other objects within the environment
- General abilities to receive, analyze and respond to input from sensors
- Communications abilities to access knowledge and data in the environment:
 - access to design data and project information data bases
 - knowledge of organizational structure and access to other agents within it
 - ability to communicate and work co-operatively with other agents
- Task planning abilities
 - understand and use design specifications for task planning
 - seek and interpret project administrative specifications for schedules, etc.
 - be able to select and locate the appropriate tools and materials
 - do the low-level planning for the task based on high-level inputs
- Extensibility and learning, the ability to assemble information from sources beyond
 - the robot, and to retain the information for the future
- Self-knowledge, to know the robot's own limits and capabilities and relate these to
 - the demands of its tasks and to requests from other agents in the environment

The sections below will examine some of these in more detail and, where appropriate, indicate current research activities that may contribute to this overall body of knowledge and abilities.

Space, Time and Mechanics

The logical starting point for a robot's model of 3-D construction could be some of the more advanced work on 3-D computer-aided design (CAD) systems now becoming the basis for graphical communications. For example, the Bechtel-Hitachi effort that has led to the software called *Construction CAE*⁴ has many of the components that enable construction processes to be simulated graphically and in scaled time against the background of realistic 3-D images of the facility being designed. It takes only a small leap of imagination to substitute real automated machines for the graphical images of construction equipment moving on the CRT screen.

The core software should also implement basic physical mechanics to monitor agent motion and location. Rapid progress in this area is being made by researchers under Professor Jean-Claude Latombe, Director of the Computer Science Robotics Laboratory at Stanford. In the past year or so they have improved motion planning algorithms to run over 100 times faster than the best previously available, and to handle over 30 degrees of freedom in three dimensions where others have struggled to handle 5 or 6 in two dimensions.^{10,19} In path planning and vehicle simulations they can manage operations like backing semi-trailer trucks through complex mazes of objects, parallel parking a car, etc., all with the automated path planner.^{3,9,20} There seems to be a great need for applying such core software work in evolving machines like the pipe-manipulator at the University of Texas,¹ whose 8 degrees of freedom seem to baffle even some of the best human operators.

Sensors

While specific sets of sensors on a given robot will be oriented toward a given application or class of applications, there are some general capabilities related to receiving and analyzing the data from the sensors that can be developed in a generic fashion and thus be included in the core intelligence. For example, machine vision and pattern recognition are fundamental problems that currently are receiving the attention of many researchers. Today this particular technology lacks fast enough processing capability to handle the types of complex images found in general construction work, but as breakthroughs in theory are made they will find their way to this field. Numerous types of simpler sensors for force, distance, temperature, etc. are already being included in construction robots. In almost all cases, however, custom programming is done to integrate them with the needs of a particular machine. As more and more machines become available, the repetition of

this custom effort should be supplanted by generic package software that can be adapted to particular needs with little or no programming. Already this has happened for laboratory data acquisition and processing (e.g., *LabView*TM for the MacintoshTM), and similar approaches could be taken for construction robots.

Language and Communication

To support this new breed of intelligent machines, issues of language and communication need to be addressed soon. Machines must evolve from isolated to cooperative states, either laterally or hierarchically. Researchers may begin by better understanding *what* information is necessary for communication—both in type and content (e.g., bitmaps, graphics, standards, knowledge, design data, task specifications, etc.)—and *how much* information to pass. Then they should address the theoretical core of languages: syntax, semantics, and pragmatics. The objectives will be to design the communication system for robotic agents working in a construction knowledge environment. Computer Science Professors Terry Winograd^{17,18} and Yoav Shoham¹⁵ are presently focusing advanced research on language issues for construction robots under two separate CIFE projects at Stanford.

One complexity in language design is that different communication tasks require different levels of knowledge representation. As an example, two robots working together might need to communicate with each other about the geometry and the exact positioning of an object. A *syntactic* representation with built-in words for geometric shapes and different types of motion may be more efficient for this robot-to-robot communication. To update a central computer's database about the status of their task, a *functional language* with descriptive codes and numeric quantities may be enough. Finally, to inform their human supervisor that a part is broken, they do not need to convey the graphical representation or the precise coordinates, just a general idea of what has happened, perhaps in a form more like *natural language*.

Another issue related to the design of a language is the level of implementation. Should a robot have enough background knowledge to infer missing parts of a discourse commonly left out in human conversation? Or should it be given complete and detailed instruction at every step? For example, if a human tells a robot driving a truck that a tire is flat, should the robot know about flat tires and thus know to stop, or should it require explicit instructions? The choice of a high-level implementation with background-knowledge may seem obvious, but, considering the dynamic nature of field project sites, providing robots with even a fraction of this knowledge would be a major undertaking, and supplying all the knowledge is impossible.

Implied in communication is some means to provide a robot with knowledge of a type of organizational hierarchy (e.g., so the robot will know how to reach a central supervisor agent—whether human or perhaps an expert system—to which task agents can turn for help). While I have listed it here under communication, it is actually part of a major research area of adapting the vast knowledge of human and organizational behavior to the more limited needs and capabilities of robot agents.

Task Planning

There are several efforts to apply advanced artificial intelligence techniques to the rudimentary planning of construction operations. Examples include site layout planning,¹⁶ multi-agent coordination,⁸ and some very interesting work to deduce plans and schedules directly from CAD drawings in their electronic form. For example, Shimizu's Mr. Kenji Ito, working for the last two years as a Visiting Fellow at CIFE, has developed CIFE CAD,⁶ which can take AutoCAD data as input and produce a schedule that recognizes the sequence of operations implied by gravity, adjacency, enclosure, etc. Such research may well lead to the ability to convert design specifications into incremental task objectives for robots.

Extensibility

Figure 2 represents a machine seeking external knowledge in spheres of increasingly difficult reachability. The nucleus of the figure shows the knowledge and data that might be pre-programmed within the robot itself and thus be immediately accessible. The next layer out could be the knowledge and data available nearby from supervisors or other agents (human or robot), or from computer databases or expert systems on-site. The third layer might still be within the robot's organization—perhaps an expert at the home office. Beyond that could be the world at large—material suppliers, equipment manufacturers, consultants and public databases. The concept of *graded reachability* in the figure, or extensibility in general, implies that, while it becomes more and more difficult and time-consuming to tap data and knowledge further and further from the robot agent's knowledge core, the agent still should have enough general knowledge and communications abilities to (a) recognize the need for knowledge beyond its own sphere, and (b) to know at least where to begin, whom or what to ask, or how to at least start to obtain the knowledge it needs to perform or continue its tasks.

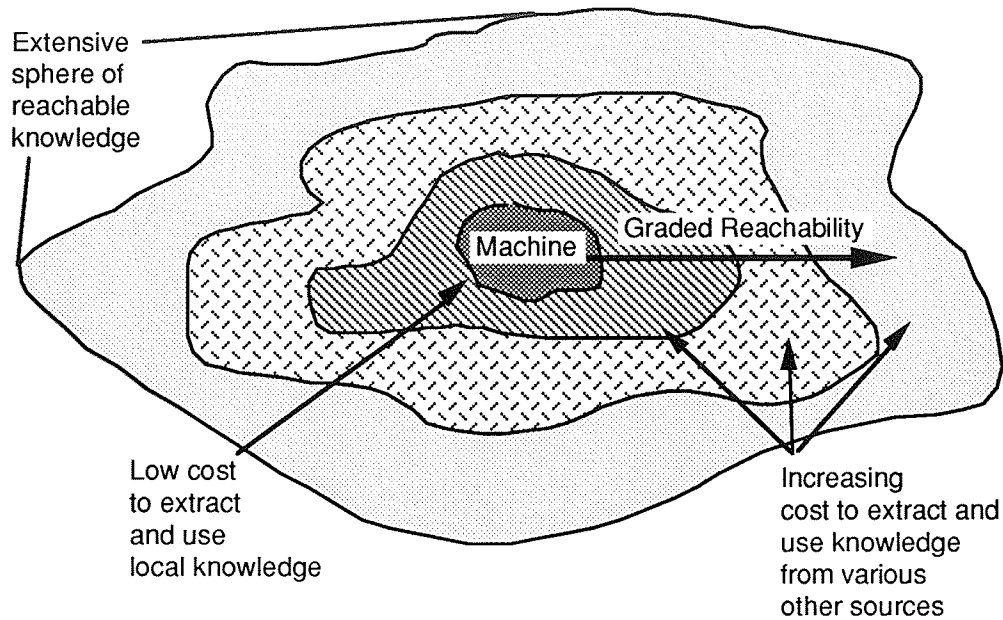


Figure 2—Accessibility of Knowledge in the Environment for Construction Automation

Self-Knowledge

To respond intelligently to unexpected situations, robots must know their own capacity and the limits thereon, if only to seek help from a human. *Self-knowledge* can help a robot determine its capability to perform tasks safely and efficiently. Self-knowledge is also essential if a robot works together with other machines or humans. Robot self-knowledge should be accessible and interpretable not only by the robot itself, but also by a central knowledge manager. The former will enable the robot to evaluate its own capability, while the latter is important when the robot needs help from the manager or wants to notify other robots about its state. Research should focus on designing a knowledge base for self-knowledge which is flexible, efficient and compatible with the general language design and communication requirements presented earlier in this paper.

Additional research is needed to define the self-knowledge needs and the environmental-knowledge interpretations of the robotic agents. The method of investigation for the environment could involve formulation of knowledge about the tasks and processes in general and about suitable tasks in particular. The self-knowledge might be divided into two parts: machine-specific self-knowledge like self-diagnostics and error-recovery; and environment-specific self-knowledge like machine weight, physical

dimensions, lifting capacity and available power. The next step would be to match suitable language constructs to the representational needs.

CONCLUSION

This has been a brief overview of a complex fabric of interwoven research needs. Machine agents in complex field environments will have to deal with many difficulties. The only tractable approach in the short term is to limit the use of machines to environments that are—or can be—carefully structured, or to leave most of the control and sensing with human agents. Researchers must look further to develop the underpinnings for more capable machines. We must discover and formulate the general computer-based reasoning and communications core for machines that would provide them with the ability to deal with unexpected events to a greater extent and exploit the opportunities provided by the knowledge environment. In particular, the kind of machines envisaged will be able to deal with uncertainty, adapt to a dynamically changing environment, be able to seek knowledge beyond their spheres, and work together in teams to perform complex tasks. Future research should provide a more robust basis for integration of machines with human organizations, with other machines and with a multitude of tools of production.

ACKNOWLEDGEMENT

For this research we are fortunate to be part of the Center for Integrated Facility Engineering (CIFE) at Stanford, which is operated jointly by the Departments of Civil Engineering and Computer Science and has corporate members from the design, construction and computer industries and representatives from owners and government agencies. We are also indebted to The National Science Foundation, which, through a series of grants (including MSM 84-06852, MSM 86-14238 and MSM 89-05573) has provided the intellectual freedom to explore the basic research that led to the formulation of the research agenda described herein.

REFERENCES

1. Alciatore, D. G., P. J. Hughes, A. E. Traver, and J. T. O'Connor, "Development of an Ergonomic Control System for Large Construction Manipulators," *Proceedings of the Sixth International Symposium on Automation and Robotics for Construction*, San Francisco, California, June 6-8, pp. 183-188.

2. Barr, Avron, and Edward A. Feigenbaum, 1981-2. *The Handbook of Artificial Intelligence*, HeruisTech Press, Stanford California, William Kaufmann, Inc., Los Altos, California, Volume 1, 1981, Volume 2, 1982, Volume 3 (by Paul R. Cohen and Edward A. Feigenbaum), 1982.
3. Barraquand, Jerome, and Jean-Claude Latombe, 1989. *On Nonholonomic Mobile Robots and Optimal Maneuvering*, Technical Report No. 015, Center for Integrated Facility Engineering, Stanford University, Stanford, California, July.
4. Ditlinger, Stephen, and Kermit Gates, 1989. "Animation/Simulation for Construction Planning and Scheduling," *Proceedings of the Sixth International Symposium on Automation and Robotics for Construction*, San Francisco, California, June 6-8, pp. 491-498.
5. Hasegawa, Yukio, 1989. "'Cutting Edge' of the State of the Art of Construction Robotization in Japan," *Proceedings of the Sixth International Symposium on Automation and Robotics for Construction*, San Francisco, California, June 6-8, pp. 27-32.
6. Ito, Kenji, Yasumasa Ueno, Raymond E. Levitt, and Adnan Darwiche, 1989. *Linking Knowledge-Based Systems to CAD Design Data with an Object-Oriented Building Product Model*, Technical Report No. 017, Center for Integrated Facility Engineering, Stanford University, Stanford, California, August.
7. Keirouz, W., Rehak, D., and Oppenheim, I., 1988. "Issues in Domain Modelling of Constructed Facilities," *Fifth International Symposium on Robotics in Construction*, Tokyo, Japan, June 6-8, pp. 341-350.
8. Koo, Charles C., 1987. *Synchronizing Plans among Intelligent Agents via Communication*, thesis submitted to Stanford University in partial fulfillment of the requirements for the degree of Ph.D., August.
9. Latombe, Jean-Claude, 1988. *Global Path Planning Approaches for Material Movements in a Worksite*, Technical Report No. 003, Center for Integrated Facility Engineering, Stanford University, Stanford, California, May.
10. Latombe, Jean-Claude, 1988. *Generation of Sensor-Based Motion Strategies in the Presence of Uncertainty*, Technical Report No. 006, Center for Integrated Facility Engineering, Stanford University, Stanford, California, June.
11. Niwa, Kiyoshi, 1990. "Human-Computer Cooperative Systems in Project Risk Management," *Proceedings of the Symposium of the Center for Integrated Facility Engineering*, Stanford University, Stanford, California, March 26-27.
12. Paulson, Boyd C., Jr., 1985. "Automated Control and Robotics for Construction," *Journal of Construction Engineering and Management*, ASCE, Vol. 111, No. 3, September, pp. 190-207.

13. Paulson, Boyd C., Jr., Thomas Froese, and Lai-Heng Chua, 1989. "Simulating the Knowledge Environment for Autonomous Construction Robot Agents," *Proceedings of the Sixth International Symposium on Automation and Robotics for Construction*, San Francisco, California, June 6-8, pp. 475-482.
14. Paulson, Boyd C., Jr., Nadeem Babar, Lai-Heng Chua and Thomas Froese, 1989. "Simulating Construction Robot Agents and their Knowledge Environment," *Journal of Computing in Civil Engineering*, ASCE, Vol. 3, No. 4, October, pp. 303-319.
15. Shoham, Yoav, 1990. *Agent-Oriented Programming*, Technical Report, Department of Computer Science, Stanford University, Stanford, California, forthcoming.
16. Tommelein, Iris D., Raymond E. Levitt, and Barbara Hayes-Roth, 1989. "SIGHTPLAN: An Artificial Intelligence Tool to Assist Construction Managers with Site Layout," *Proceedings of the Sixth International Symposium on Automation and Robotics for Construction*, San Francisco, California, June 6-8, pp. 340-347.
17. Winograd, Terry, 1983. *Language as a Cognitive Process*, Addison-Wesley, Menlo Park, California.
18. Winograd, Terry, 1989. *Language/Action Approach to the Computer Integration of Construction Work*, Seed Research Proposal to the Center for Integrated Facility Engineering, Stanford University, Stanford, California, December.
19. Zhu, David, and Jean-Claude Latombe, 1989. *Contingency-Tolerant Robot Motion Planning and Control*, Technical Report No. 014, Center for Integrated Facility Engineering, Stanford University, Stanford, California, July.
20. Zhu, David, and Jean-Claude Latombe, 1989. *New Heuristic Algorithms for Efficient Hierarchical Path Planning*, Technical Report No. 018, Center for Integrated Facility Engineering, Stanford University, Stanford, California, August.

