

SIMULATION-BASED REAL-TIME DECISION MAKING FOR ROUTE PLANNING

Jin Joo Lee
Paul A. Fishwick

Department of Computer & Information Science & Engineering
University of Florida
Gainesville, Florida 32611, U.S.A.

ABSTRACT

Decision making is an active area of research in simulation, systems engineering and artificial intelligence. One subset area of decision making, automated route planning, is covered in our paper with our approach being based on the technique of simulation rather than on purely heuristic or geometric techniques. This new technique is called simulation-based planning. Simulation-based planning is useful for route planning under various conditions including uncertain locations and events with potential adversarial activity. We claim that it is only by using simulation that one can make the most effective plan in uncertain and complex environments. An experimental design is outlined along with our plans for further development.

1 INTRODUCTION

At a fundamental level, general AI planning, decision making, intelligent control and route (path) planning in robotics, all strive to solve a common problem—based on some model of a given process, determining what actions will affect the process in a desired way. The problem appears different because each area deals with different levels of abstraction and applications.

Planning becomes very complex for real world planning problems that take place in an environment over which the planner has no control, such as with another agent or an enemy or when there is uncertainty of available information or agents' reactions. In such cases, accurate prediction of the resulting states of plan execution will be difficult. To overcome this increase in complexity of reasoning, many new approaches have been introduced (Schoppers 1987; Dean and Kanazawa 1987; Kanazawa and Dean 1989; Hammond 1989). To handle these uncertainties in real time is the ultimate goal of any planner.

Simulation-Based Planning (SBP) can solve the problem of prediction of uncertain environment by allowing the use of multiple simulations of simulation models to predict the behavior of individual objects using data sampled from a predetermined distribution. Moreover, SBP can produce plans in real time since we can allow simulation of plans at different levels of abstraction. A high-level, simulation can be done by simply sampling numbers from a distribution or a complex low-level simulation can be done by simulating the state change in greater detail at each time step. Our previous work in military mission planning with Computer Generated Forces shows preliminary results of our SBP approach (Lee and Fishwick 1994).

The related background research areas to planning are discussed in Section 2. In Section 3, we discuss SBP as a general methodology. Our example problem—the rover route planning problem—is presented in Section 4 and experimental design issues are discussed in Section 5. Finally, conclusions appear in Section 6 and future work appears in Section 7.

2 BACKGROUND RESEARCH AREAS RELATED TO PLANNING

To view the problem of route planning as part of a larger picture and to gain better insight into the nature of our problem, we will overview different areas that are relevant to planning.

2.1 Artificial Intelligence

The general problem of planning in AI is commonly identified with problems that are highly conceptual where actions are of the form “Go To Supermarket” and “Buy milk” (Russell and Norvig 1995). A plan is an ordered set of these high-level actions. Here, the concern is not how one will physically (at a detailed level) get to a supermarket, but rather on the ordered set of actions whose logical effects will satisfy the goal. STRIPS (Fikes, Hart, and Nilsson 1972a;

Fikes, Hart, and Nilsson 1972b) is a classical example of such an approach to planning. This approach is reasonable if the execution of the produced plan is not a responsibility of the planner. Difficulties arise when execution becomes part of the planning system.

2.2 Intelligent Control

Intelligent control deals with problems that are more physical and less conceptual. The problem of steering a cargo ship to a desired heading (Antsaklis and Passino 1993) is a typical problem in intelligent control. Even though this whole task can be just simply stated as “steer the ship to heading x”; at the control level, we are more concerned with tuning the control input to physically steer the ship to a precise heading.

2.3 Decision Science and Game Theory

Decision science involves the creation of decisions based on a game-theoretic foundation. Given the current “state of the world,” one can embark upon several courses of action (decisions) each of which will yield a payoff or utility (Luce and Raiffa 1957). Games can be naturally extended to continuous systems (Basar 1995)(often found in simulation) by equating the input (or forcing) function to a continuously changing *decision* which alters the payoff given the corresponding state changes.

3 SIMULATION-BASED PLANNING

In the simulation literature, simulation is defined as “the discipline of designing a model of an actual or theoretical physical system, executing the model on a digital computer, and analyzing the execution output” (Fishwick 1995). In the planning literature, planning is defined as using models to formulate sequences of actions and given a sequence of actions, models are used to simulate the future as it would occur if the actions were carried out (Dean and Wellman 1991). So simulation provides the robot with information which can be used to suggest modifications or to compare the proposed sequence with an alternative sequence. Humans, who plan using a similar overall approach, have models built and stored in their brains for most objects or systems. These models are used to formulate sequences of actions which would occur in the future if a plan was executed. Therefore, once simulation models have been built for a system, simulation can be used as a tool to provide the system with information useful for evaluating an hypothesis. It is logical that we employ simulation within the planning process to gather information

about each candidate plan (sequence of actions) and to compare them. Some recent work (West, Mellon, Ramsey, Cleary, and Hofmann 1995) presented in this conference also bears significant relation to our work. Their work focuses on speeding up the execution of models for military strike-planning using parallel and distributed simulation. Our approach can potentially use their method for further increasing the speed of plan and route evaluation, in addition to using parallel replication by simulating, in parallel, models with separate factor values.

Once a plan is chosen for execution, the simulation data that was generated during the planning process can be used to match with the current real world state. This can be compared to a common technique used in adaptive control theory where a reference model is compared with the actual performance data in order to tune the controller to a desired state (Antsaklis and Passino 1993).

4 EXAMPLE PROBLEM: ROVER ROUTE PLANNING

Our focus is near-optimal route planning. Route planning is in-between the higher level of symbolic AI planning and the lower level of intelligent control. There are several application areas that are related to route planning. Mission planning within the military domain almost always involves route planning. Routes can greatly affect the success of the whole mission, whether the mission takes place on ground or in the air. Some of our previous research was focused on this aspect; we used simulation-based planning to perform mission planning at the company level of command (Lee and Fishwick 1994). Using simulation-based planning in the military domain was a natural extension of the already existing wargame approach where the mission plans were tested off-line via constructive simulations. Related work by (Czigler, Downes-Martin, and Panagos 1994) demonstrates the usefulness of simulation as a decision tool.

Robot route (i.e. path) planning is another application area. If there is little uncertainty involved, as is sometimes the case in many robot route planning problems, the existing approaches such as potential fields (Barraquand, Langlois, and Latombe 1992) do quite well. But, when uncertainties exist in the environment, these methods alone cannot produce good results. The main algorithm of producing a graph of traversable paths and searching the graph for a best route is still the same, however. The part that is different is how we measure the *goodness* of a route. If the goal is to select a route that is the shortest in distance, we can use any of the standard algorithms

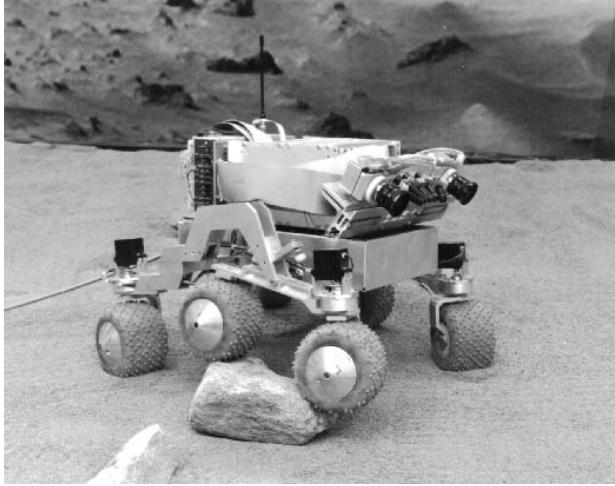


Figure 1: Mars Microrover. Provided by permission of Jet Propulsion Laboratory (Matthies, Gat, Harrison, Wilcox, R., and Litwin 1995)

that exist for finding shortest paths in a graph. But, if the problem is in an environment that is unknown or uncertain, we must use different ways to evaluate each path. Our claim is that we can use simulation to quantitatively compute the outcome of a future action. The way we do this is by using simulation models for each entity or object in the world using data sampled from a probability distribution, and performing multiple simulations to obtain an estimate of the state variables.

In 1996, NASA plans to launch a spacecraft to Mars to explore the environment of the planet (Matthies, Gat, Harrison, Wilcox, R., and Litwin 1995). The spacecraft will carry an 11 kg rover, called the Microrover, that will move around the vicinity of the landing site to explore the territory for a duration of approximately 1 to 4 weeks. Figure 1 shows the Microrover traveling over a rock.

Because the Martian surface is not completely known, JPL is undergoing a process of performance evaluation of the rover's autonomous navigation system with varying terrain characteristics. The Microrover testbed contains the Microrover vehicle and an indoor test arena with overhead cameras for automatic, real-time tracking of the true rover position and heading. In the arena, they have created Mars analog terrains by randomly distributing rocks according to an exponential model of Mars rock size and frequency created from Viking lander imagery. Figure 2 shows a *nominal* Mars terrain which was adopted from (Matthies, Gat, Harrison, Wilcox, R., and Litwin 1995). The term *nominal* is explained in more detail in Section 5. JPL has decomposed the

rover navigation task into four functions: 1) goal designation; 2) rover localization; 3) hazard detection; and 4) path selection. Although these four functions are integrated, for purposes of research, we will focus mainly on route selection. The first three functions are largely connected with problems in sensors (stereo camera pair, wheel encoders, solid state turn rate sensor, light stripe ranging sensor, inclinometers) and their uncertainties. Path selection uses the information gathered from these three functions to navigate to a goal destination. Currently the route selection is achieved by a simple behavior control algorithm which is reactive and does not take any excess knowledge—such as maps—into account. Our approach is to use simulation at this stage of route planning to select the route to the goal. Once a more detailed map is constructed of the vicinity, the planning will involve multiple simulations of each virtual route to the goal.

4.1 Simulation-Based Route Planning System

Figure 4 illustrates the basic components of our Simulation-Based Route Planning system. Initially, the planner takes the goal location of the route as input and selects a route plan for output. This selected plan is the input to the Control Subsystem which performs a supervisory control of the process. The output of the rover process is the actual sensory output of the rover. The sensory output will include camera images, hazard detections and position information. Along with the plan, a simulation log (the simulation data that was produced previously during the plan evaluation process) of the chosen plan is provided as input to the Control Subsystem. This can be used to serve as a reference model to track the state of the execution in order to monitor its progress towards the goal. The monitoring information can be used further to tune the system towards the goal (i.e. correct its route or position estimation) or to generate a failure signal to the planner as soon as it decides that the current route is unlikely to succeed.

4.1.1 Planner

The planner has three major modules:

1. The **Route Generator** uses a map of the vicinity leading to the goal to extract rocks that are large enough to be considered as obstacles. The threshold size is determined by the size and abilities of the rover. In the case of the Mars Microrover, it has been determined to be rocks that have diameters larger than 23 cm. In Figures 2

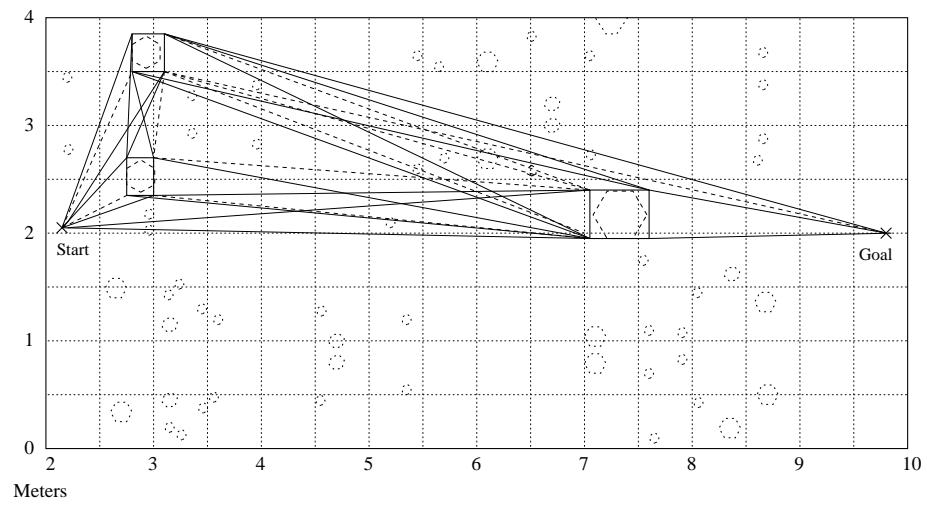


Figure 2: Map of Nominal Mars Terrain with Visible Routes

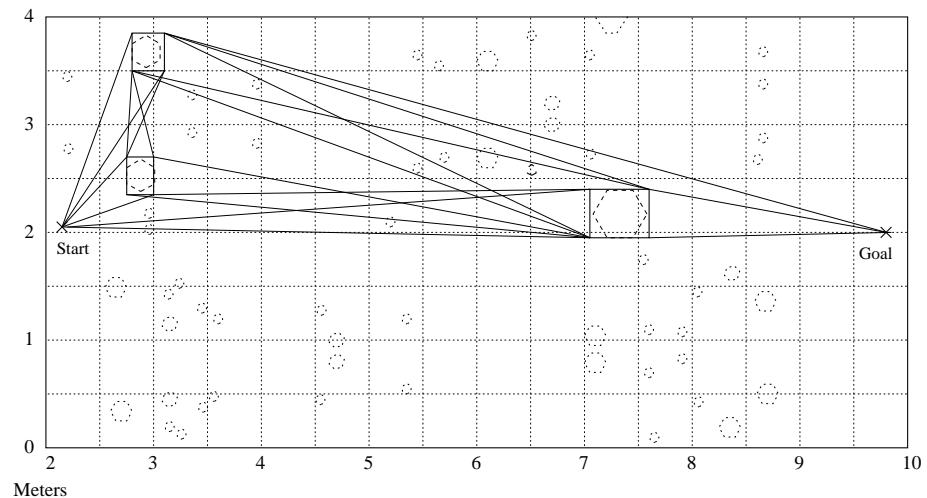


Figure 3: Map of Nominal Mars Terrain with Reduced Visible Routes

- and 3, the rocks are represented by hexagons drawn with dotted lines. There are three rocks which are big enough to be considered as obstacles. A bounding box is drawn around each of the rocks. A visibility graph is produced connecting the start location and the goal location via these bounding boxes (Figure 2). The dotted lines in the figure represent routes that will be eliminated via the *reduced* visibility graph method. The visibility graph is a common approach used to create paths between obstacles (Latombe 1991) by connecting vertices of each obstacle to every other obstacle, including the start and goal location. The reduced visibility graph is a graph having the minimal Euclidean route length between the start and goal location and the obstacles. This is what appears in Figure 3. The reason for generating routes between obstacles is because we use the obstacles as landmarks during navigation. If a direct (as opposed to indirect which means the route touches an obstacle) route exists to the goal location, the reduced visibility graph will produce it since this *is* the shortest route to the goal.
2. The **Simulator** simulates each route of the produced graph and records the result. The simulation is based on the physical and empirical models of the rover and the terrain. The physical model includes specific characteristics such as translation (max 0.67 cm/sec) and rotation speed, step climbing ability (max 19.5 cm), and specification for the hazard detection sensors (range of view is about 120 deg. with 30 cm max distance). We also incorporate empirical models for dead-reckoning error and hazard detection error (failure rate of 1 in 1000 given that hazard frequency is 1/100). For the terrain, we use the Moore's model of rock size-frequency distribution to create terrains of different rock sizes and frequencies. More detail about this is discussed in Section 5. Using these models of the rover and the Martian surface, we perform multiple simulations of each route. There are several ways to proceed in the simulation. To perform in real-time, it's preferable that we proceed in a breath-first manner; for example, in Figure 2, we simulate routes to the first set of obstacles and then to the second set and so on. This approach is similar to the approach we have taken in our previous work in mission planning (Lee and Fishwick 1994). To reduce the amount of computation, we use the A* search method if we can build a heuristic function which can estimate the

cost of the remaining route. Another possibility is the branch and bound method used in the area of Operations Research. Due to uncertainties which exist in the models, simulations must be performed multiple times using the available stochastic information to reduce the variance of the outcome variables.

3. The **Plan Evaluator/Selector** evaluates the results of the simulations and selects a route for execution. Currently, three elements are considered: 1) the outcome of plan failure or success; 2) the total time elapsed; and 3) final position of the rover. These elements can be combined into a single score but, because the objectives may be different in different situations, we evaluate them individually and then select a plan according to a predetermined criterion.

5 EXPERIMENTAL DESIGN

In simulation, experimental design is a method of choosing which configurations (parameter values) to simulate so that the desired information can be acquired with the least amount of simulating (Law and Kelton 1991). In experimental design terminology, the input parameters and structural assumptions composing a model are called factors and the output performance measures are called responses. Our experimental design approach to the SBP method is to choose different rock distributions (both in terms of size and frequency) and different routes as our factors, and the simulation results according to the evaluation function as our responses. Thus, we vary two elements of simulation within the planning process: 1) the route; and 2) the terrain. If the planning experiment is performed using the full factorial design approach, the number of combinations to simulate will be prohibitive. However, optimization techniques such as response surfaces and metamodels can be used to alleviate the problem. Currently, JPL is experimenting their Microrover by performing test runs on Mars analog terrains which has been created artificially by randomly placing rocks according to a rock size-frequency model developed by Moore (Moore and Jakosky 1989). Moore's model is based on data obtained from images taken by Viking Lander 2. Since a similar rock density is expected for the Microrover experiment, the same model can be used. The original Moore's model for rocks down to a diameter of 0.14 m is represented by $N = 0.013D^{-2.66}$, where N is the cumulative frequency of rocks per square meter with diameters of D and larger. This model predicts that about 18.8% of the landing site area is covered

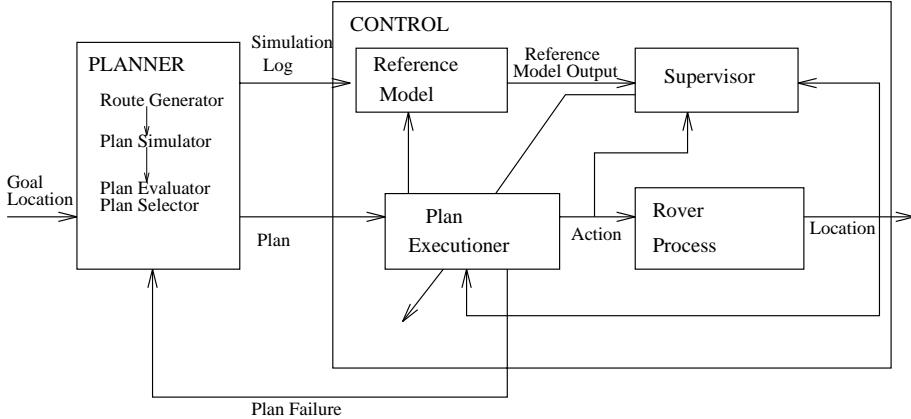


Figure 4: Simulation-Based Route Planning System

by rocks. However, the model used by JPL so far in creating the actual test terrains is based on the modal value of the surface rock cover over the whole planet, which is estimated to be at 6%. The terrain created from this model is called *nominal*. A computer simulation is under development at JPL in order to test terrains with rock frequencies ranging up to 19%.

We base our simulated terrains on this model in creating simulated Mars analog terrains. There are three major factors to consider in creating the terrain: 1) rock sizes; 2) rock frequencies; and 3) rock placements. For each simulation, the ranges of rock sizes and frequency are sampled from the Moore's model. Once the rover is on Mars, the map that is generated from photographs will be used to place rocks that are large enough to be obstacles in their absolute locations. The remaining terrain in between the large rocks can be estimated through the Moore's model during simulation. We can use any additional information that is available (distribution of rock sizes and frequency) of the landing site to design a non-uniform sampling distribution. Initially, the sampling distributions are uniform. By using visual information of the landing site, we can build rock distributions that are similar to the actual terrain characteristics. The placement of these smaller (less than 23cm in diameter) rocks will be random. If time permits, multiple simulations with different small rock placements will be performed. On a higher level, the percentage of rock cover can be varied to be between 6% to 19%. Figure 5 illustrates the variable factors in the simulation.

For the simulation, we use discrete time step simulation and the following algorithm:

```

While (Goal is not reached) do
    Sample sensor data
    Execute control action on rover

```

```

Update rover state variables
Update current clock time by ΔT
End While

```

Until the rover's state variables indicate it has reached the goal location, the planner continues above loop. Sampling sensor data involves sensing the rocks placed and sized on the map according to the previously discussed method. Depending on the size and location of a rock, the planner may also invoke the hazard detection sensors—indicating that the rover will have to maneuver around the obstacle. Smaller rocks (with diameter less than 23 cm) are ignored by the hazard detection sensor but is likely to cause the dead-reckoning error to increase. This effect can be captured by incorporating a dead-reckoning error model into the simulation. This is planned for later experiments. The control algorithm of the rover is very simple: if an obstacle is detected then the rover turns in place until the hazard is no longer detected. If there is no obstacle then it continues to move forward while turning towards the goal location. The control action is calculated to take place for ΔT seconds and the state variables (location x,y) are updated accordingly using control dynamics of the rover.

6 CONCLUSIONS

We have shown how we are able to build a real-time route planning system, the Simulation-Based Route Planning System, by integrating ideas from different but related fields of research.

There are several advantages to the method of simulation-based planning:

1. By employing models to simulate and predict

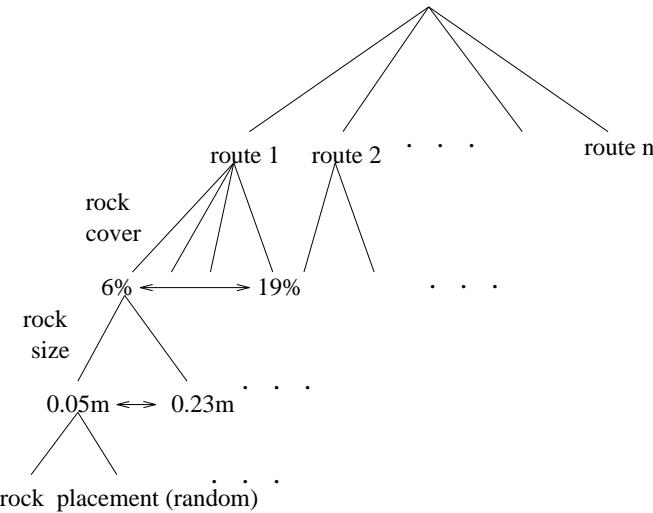


Figure 5: Sampling Range and Distribution of Simulation Factors

- the outcome of the process, we are able to capture the effects more accurately and completely (given that the models are built appropriately).
2. The use of simulation models allows us to use standard simulation analysis methods in tuning the simulation models to closely reflect actual processes.
 3. The ability to use distributed object-oriented concept in planning without having to reason about the combined effects of agent's actions or changes in the environment within one central node enables the evaluation of a plan as a natural result of simulating different models in the system.
 4. The ability to "track" progress and performance during execution allows finer tuning of the execution process.
 5. We can easily extend the set of models to include additional properties (especially properties that may be difficult to create physically such as creating the atmosphere of Mars) in testing plans.

Some potential difficulties exist in using simulation in the planning process. Defining appropriate models may be difficult and time consuming. The simulation process itself can be computationally intensive. However, we plan to overcome this problem by varying the level of aggregation of our simulation models.

7 FUTURE WORK

Once the implementation of the rover problem domain in our Simulation-Based Route Planning

(SBRP) system is finished, we will experiment with our method by building a response surface of the problem and then using various ways to search for the near optimal solutions. In the long term, we plan to extend our SBRP system to the domain of Air Force mission planning problems.

ACKNOWLEDGMENTS

We would like to thank Science Applications International Corporation (SAIC, Orlando, Florida) under contract #4515164 for partial support of this research. Continued partial support of this work will be possible through contract #4514169-12 from Rome Laboratory, Griffiss Air Force Base, New York. We would also like to thank the JPL Mars Microrover group for their permission to use Figure 1 and various data appearing in (Matthies, Gat, Harrison, Wilcox, R., and Litwin 1995). Disclaimer: this work is not funded by JPL. The authors have used Figure 1 and other rover related data as a sample application of the simulation-based planning procedure.

REFERENCES

- Antsaklis, P. J. and K. M. Passino. 1993. *An Introduction to Intelligent and Autonomous Control*. Kluwer Academic Publishers.
- Barraquand, J., B. Langlois, and J. Latombe. 1992. Numerical Potential Field Techniques for Robot Path Planning. *IEEE Transactions on Systems, Man, and Cybernetics* 22(2), 224-241.
- Basar, T. 1995. *Dynamic Noncooperative Game Theory*. Academic Press.

- Czigler, M., S. Downes-Martin, and D. Panagos. 1994. Fast Futures Contingency Simulation: A "What If" Tool for Exploring Alternative Plans. In *Proceedings of the 1994 SCS Simulation MultiConference*, San Diego, CA.
- Dean, T. L. and K. Kanazawa. 1987. Persistence and probabilistic inference. Technical Report CS-87-23, Department of Computer Science, Brown University.
- Dean, T. L. and M. P. Wellman. 1991. *Planning and Control*. Morgan Kaufmann.
- Fikes, R. E., P. E. Hart, and N. J. Nilsson. 1972a. Learning and Executing Generalized Robot Plans. *Artificial Intelligence* 3.
- Fikes, R. E., P. E. Hart, and N. J. Nilsson. 1972b. Some New Directions in Robot Problem Solving. In *Machine Intelligence 7*. Edinburgh University Press.
- Fishwick, P. A. 1995. *Simulation Model Design and Execution: Building Digital Worlds*. Prentice-Hall.
- Hammond, K. 1989. Cased-based planning. In *Perspectives in Artificial Intelligence*, Volume 1. Academic Press.
- Kanazawa, K. and T. L. Dean. 1989. A Model for Projection and Action. In *Proceedings of IJCAI-89*, Detroit, MI, pp. 985-999.
- Latombe, J. 1991. *Robot Motion Planning*. Kluwer Academic Publishers.
- Law, A. M. and W. D. Kelton. 1991. *Simulation Modeling and Analysis*. McGraw-Hill.
- Lee, J. J. and P. A. Fishwick. 1994. Real-Time Simulation-Based Planning for Computer Generated Force Simulation. *Simulation*, 299-315.
- Luce, R. D. and H. Raiffa. 1957. *Games and Decisions*. John Wiley and Sons. Later re-printed by Dover, 1985.
- Matthies, L., E. Gat, R. Harrison, B. Wilcox, V. R., and T. Litwin. 1995. Mars microrover navigation: Performance evaluation and enhancement. Technical report, Jet Propulsion Laboratory, Pasadena, CA. To appear in Autonomous Robots Journal.
- Moore, H. J. and B. M. Jakosky. 1989. Viking Landing Sites, Remote-Sensing Observations, and Physical Properties of Martian Surface Materials. *International Journal of Solar System Studies*, 164-184.
- Russell, S. J. and P. Norvig. 1995. *Artificial Intelligence A Modern Approach*. Prentice-Hall.
- Schoppers, M. 1987. Universal Plans for Reactive Robots in Unpredictable Domains. In *Int. Joint Conference on Artificial Intelligence*.
- West, D., L. Mellon, J. Ramsey, J. Cleary, and J. Hofmann. 1995. Infrastructure for Rapid Execution of Strike-Planning Systems. In *Proceedings of the 1995 Winter Simulation Conference*, Washington, D.C.

AUTHOR BIOGRAPHIES

JIN JOO LEE received the B.S. degree in Computer Science from Ewha Womans University, Korea in 1988 and the M.S. degree in Computer Science from Brown University in 1991. After receiving the M.S. degree, she was a research engineer at Human Computers Inc., Korea until 1992. She is currently a PhD student in the Computer and Information Science and Engineering department at the University of Florida. Her research interests are in AI planning, simulation and control. Lee's homepage is <http://www.cis.ufl.edu/~j11>.

PAUL A. FISHWICK is an Associate Professor in the Department of Computer and Information Science and Engineering at the University of Florida. He received the BS in Mathematics from the Pennsylvania State University, MS in Applied Science from the College of William and Mary, and PhD in Computer and Information Science from the University of Pennsylvania in 1986. He also has six years of industrial/government production and research experience working at Newport News Shipbuilding and Dry Dock Co. (doing CAD/CAM parts definition research) and at NASA Langley Research Center (studying engineering data base models for structural engineering). His research interests are in computer simulation modeling and analysis methods for complex systems. He is a senior member of the IEEE and the Society for Computer Simulation. He is also a member of the IEEE Society for Systems, Man and Cybernetics, ACM and AAAI. Dr. Fishwick was chairman of the IEEE Computer Society technical committee on simulation (TCSIM) for two years (1988-1990) and he is on the editorial boards of several journals including the ACM Transactions on Modeling and Computer Simulation, IEEE Transactions on Systems, Man and Cybernetics, The Transactions of the Society for Computer Simulation, International Journal of Computer Simulation, and the Journal of Systems Engineering. Fishwick's home page is <http://www.cis.ufl.edu/~fishwick>.