On the Role of Simulation in the Study of Autonomous Mobile Robots

Erann Gat
Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109
gat@jpl.nasa.gov

ABSTRACT

This is the third in a series of essay about various aspects of research in autonomous mobile robots. This essay examines the role of simulation. I argue that simulations should be considered scientific models rather than surrogates for real robots.

1. Introduction

Working with real robots in real environments is usually expensive and time-consuming. As a result many researchers rely on simulations. This practice has been criticized because, as Rod Brooks puts it, "Simulations are doomed to succeed." Simulations may not (and often do not) adequately capture those aspects of reality that make controlling autonomous agents in the real world a difficult problem.

The idea that simulation results should not be blindly trusted is now generally accepted, but researchers are understandably reluctant to simply abandon simulations because they are so convenient to work with. A sort of compromise position has emerged that simulators may be properly used to more efficiently search the design space of control mechanisms, but that the results should be verified on a real robot before betting your life savings on them.

This compromise has been defended on the grounds that if an algorithm doesn't work in simulation it won't work in reality, and thus simulations can provide a cost-effective filter of unprofitable methodologies. Unfortunately, it isn't true. There are algorithms that work in reality which will produce endless futile loops in a naive simulation. Furthermore, the original criticism of simulations still applies, and so it is not at all clear that the filter that simulations are supposed to provide will be all that effective.

If your goal is to do science (as opposed to simply building robots that work) then the situation is even worse. If your simulation works and your real robot fails (or vice versa), what does that say about your theory? It is tempting in such cases to adopt the view that reality never lies, but what if the real robot failed because a wire came loose? There is currently no satisfactory way to resolve such conflicts.

I will argue for adopting a completely different point of view: that simulations should be viewed as *models* rather than as *systems under study*. I will argue that this view can eliminate much of the confusion and debate, and advance both the causes of science and engineering in the field. I will begin with a general discussion of models and systems.

2. Models and systems

Science is the business of seeking truth through experimentation. This simple idea has been developed into a powerful and complex intellectual framework. The power of science is manifested in its ability to enable us to predict certain aspects of future events. This is done through the use of scientific *models*, intellectual constructs that can be manipulated to generate predictions. Newtonian physics, for example, is a scientific model that generated a prediction that comet Shoemaker-Levy would collide with Jupiter in 1994.

The intellectual apparatus that has been developed to support the generation, manipulation, and communication of models is so sophisticated that it is often mistaken for the whole of science. But science isn't consummated until one does an experiment to test a prediction generated by a model. The mathematics underlying Murray Gell-Mann's theoretical derivation of the existence of quarks is a canonical example of a scientific triumph, but it didn't become science until someone fired up the atom smasher and collected some data.

You can't do science about robots without firing up a robot. This is inescapable; to do an experiment, which is to do science, you have to interact with the system you are studying, otherwise you are counting angels on the head of a pin. The unfortunate reality is that, given the

technological state of the art, to do an experiment with a real robot involves a substantial amount of beating your head against the wall just trying to keep the robot from beating its head against the wall. After spending hours or months on menial chores like fixing broken wires and trying to figure out why the radio modem isn't working again, it is extremely tempting to use a simulation as a surrogate. The result may be convenient, but it isn't science.

Fortunately, this does not mean that simulations are useless. In the next section I propose an alternative.

3. Simulators as models

Let us take a closer look at the way in which science often appears to be getting done in mobile robots, but really isn't. A researcher proposes a theory, i.e. a model. This model usually takes the form of a set of mathematical equations. The researcher generates a set of (often tacit) predictions, and then presents results of "simulator experiments" to show that the predictions are borne out. The simulator is playing the role of the *system under study*, substituting for the inconvenient real robot. This is the current, unscientific state of affairs.

There is a procedure which is sometimes used to bolster the credibility of a simulator as a system surrogate. An experiment is done on a real robot, and then the same experiment is repeated on the simulator to show that the simulator produces the same result as the real robot (e.g. [Gat92a]). This is a step in the right direction, but the utility of this approach is limited. When the simulator is used to produce results under different circumstances than the ones it was verified under it can no longer be trusted. The only way to really be sure that your results are real is to verify the simulator under all circumstances of interest, and then you are back to doing all your experiments on a real robot anyway, and the simulator is moot.

But suppose that we reversed the order: rather than run an experiment on a real robot and then reproducing it on a simulator, we first run an experiment on the simulator and then reproduce the results on a real robot. Now we are doing science! Why? Because the simulator is playing the role of a model.: it is generating predictions which are then verified with real experiments.

There are many fortuitous consequences of adopting this view. For one thing, we can dispense with the classical analytical models of robots, which are cumbersome, and appear to be woefully inadequate for describing the behavior of robots [Gat92b]. In their place we now have computational models (simulations), which are still mathematical models, but more complex (though easier to deal with) and thus better able to capture the complex behaviors observed in real robots.

A second advantage of this approach is that it does not require much to change in the day-to-day business of doing robotics research. It is more a change in one's way of thinking. To oversimplify things a bit, what I am advocating is simply that the thought running through your mind when you run a simulation change from, "I am doing an experiment," to, "I am generating a prediction from a model."

A third advantage is that it allows the incremental and collaborative construction of very complex scientific models. Because models under this view are computer programs, we can apply software engineering techniques to their construction. In the next section I describe a simulation designed for this purpose.

Finally, adopting this view will enable us to defend our work as real science rather than the alchemy that it currently appears to be. An interesting side-effect of my approach is that it allows us to claim to be making progress even when things don't work! In science, the refutation of a model by experimental evidence is just as important (sometimes even more important) than a confirmation.

4. An example

JPL is currently conducting a study of Mars rover navigation [Gat95]. The design of this study is based on the simulations-as-models idea described above.

In terms of the focus of this paper, the centerpiece of the study is a simulation called EROS (Erann's RObot Simulator). EROS is designed to be a modular, extensible simulation. It is aggressively object-oriented; everything in EROS is an object. Robots are constructed by "installing" sensor and actuator objects into robot chassis objects. Completed robots are then installed into world objects, along with non-robot objects such as obstacles. The simulation naturally supports multiple robots by simply installing more than one robot into a world, and multiple simulations by simply making more than one world.

The lowest level functionality provided by EROS is the efficient detection of intersections among the physical boundaries of world objects. EROS currently uses a two-dimensional polygonal model for efficiency, but this (like everything else in the system) is easy to change.

The other component of the study is a robot called Rocky 3.2, which has the same size, chassis design, computer and sensor suite as the robot that NASA plans to send to Mars in 1996. Experiments with this robot are conducted in a sandbox instrumented with four overhead cameras that track the rover's position and orientation as it moves.

A model of the rover has been implemented within the EROS framework. This model has successfully post-dicted the results of an initial series of experiments performed on the real robot. We are currently in the process of generating our first quantitative predictions in a new operating regime. We hope to confirm these predictions with experiments by the end of this summer.

One unique feature of this work is that our predictions are statistical; we are predicting the probability distributions of certain performance metrics. It is probably impossible to predict the actual behavior of the robot in any particular situation because the performance is chaotic (in the mathematical sense) and two runs with apparently identical initial conditions can give radically different results. In such circumstances statistical prediction is the best we can do.

5. Summary

I have argued for using simulations as scientific models rather than as experimental surrogates for real robots. Running a simulation is not "doing an experiment" but rather "generating a prediction." When such a prediction is subsequently verified (or refuted) by an experiment on a real robot, the result is consummated science.

We are implementing this view in a real study of navigation performance in rough terrain. We have completed and calibrated a simulation, generated predictions, and are currently conducting experiments to confirm those predictions. Our predictions are statistical; the chaotic nature of the behavior of mobile robots make case-by-case prediction impossible.

References

[Gat92a] Erann Gat, Integrating Planning and Reaction in a Heterogeneous Asynchronous Architecture for Controlling Mobile Robots, *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI)*, 1992.

[Gat92b] Erann Gat, On the Role of Theory in the Control of Autonomous Mobile Robots, AAAI Fall Symposium on Applications of AI Theory to Real-world Autonomous Mobile Robots, 1992.

[Gat95] Erann Gat, Towards Principled Experimental Study of Autonomous Mobile Robots I, in preparation.

Addendum

Following are my responses to some of the specific questions raised by the symposium chairs.

<u>Coordination -- How should the agent arbitrate/coordinate/coperate its behaviors and actions? Is there a need for central behavior coordination?</u>

Yes, there must be centralized coordination. Attempting to arbitrate multiple conflicting commands is a fundamental mistake because there is not enough information provided to the arbiter to make the right choice. However, the centralized coordinator should not control the robot's actuators directly. Instead, it should provide indirect control by selecting among alternative transfer functions in a reactive layer.

<u>Interfaces-- How can human expertise be easily brought into an agent's decisions?</u>

By selecting an representation ontology that matches the one that humans use to instruct each other. I favor an ontology based on Agre and Chapman's plans-as-communications model.

<u>Representation-- How much internal representation of knowledge and skills is needed?</u>

I think this is the wrong question. How much you *need* depends on what you're trying to do. Many practical tasks can be done with very little representation. The right question is: is there a substantial benefit to having more than the minimal amount? I think the answer is yes. There appears to be a practical upper limit to

what can be done using a hand-engineered approach like subsumption. To go beyond this limit explicit representation is necessary in order to manage the complexity.

What is gained by using a multi-level architecture versus a monolithic architecture?

Multi-level heterogeneous architectures allow different components of the system to be optimized for different tasks.

How much does each component of an agent architecture have to know about the other?

To quote Matt Mason: the right amount. (i.e. there is no hard and fast answer. It depends on what the component does.)

Performance-- What types of performance goals and metrics can be used?

I think this is the most important question in the list, and one that is very difficult to answer. Let me contrast my answer with Jim Albus's answer, quoted here:

> Success is the ultimate metric for intelligent agents. Do they accomplish their given goals? How well? How fast? How many errors along the way?

This is right, but doesn't go nearly far enough.

We need objective performance metrics if we are to do science. Unless we can measure our success we cannot know if we have succeeded, and we cannot compare alternative approaches. The problem is that performance can be measured along a huge number of incommensurate axes: speed, cost, failure rate, severity of failures, extent of the conditions under which failures occur, flexibility, etc. etc. etc. This is a Hard Problem, one that deserves more attention than it has gotten.

Simulation -- What, if any, role can advanced simulation technology play in developing and verifying modules and/or systems?

This is the question I address in the main body of my submission. To summarize briefly: Simulations should be considered mathematical models — very complex and sophisticated ones. They should play the same role as any other mathematical model: to enable us to make predictions that are subsequently verified or refuted by experiment. Too often in robotics simulations are used as surrogates for the systems under study. (You see this every time you see an analytical model "verified" by "experiments" performed on a simulator.) This is a mistake and has impeded progress. Unfortunately, the result has been a vilification of the simulators themselves rather than the manner in which they are used.

Can we have standard virtual components/test environments that everybody trusts and can play a role in comparing systems to each other?

Absolutely. We can, we should, and we must. EROS is an attempt to provide a common framework for implementing such test environments.

How far can development of modules profitably proceed before they should be grounded in a working system?

No one knows. In general, the more reality checks the better, but reality checks are expensive. Only time and experience will tell us where the optimal tradeoff point is. This issue should be resolved on pragmatic grounds rather than dogmatic ones.