

THE SCIENTIFIC RELEVANCE OF ROBOTICS

*Remarks at the Dedication of the
CMU Robotics Institute
December 9, 1980*

Allen Newell
Department of Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

Introduction

I am absolutely delighted to be able to join in this morning to offer my reflections on the occasion of the official beginning of the Robotics Institute. Beginnings are full of promise and potential. This one is no exception. What the Robotics Institute will become -- what effects it will have, both witting and unwitting -- are for the future to tell. What we all have now is a sense of adventure and anticipation.

The appearance at this dedication ceremony of Admiral Baciocco, from the Office of Naval Research, Mr. Murrin, from the Westinghouse Corporation and myself, from Carnegie-Mellon University, is significant. That significance can be viewed narrowly, of course. We represent the three particular organizations that are most responsible for bringing this new endeavor into being. But viewed more broadly, the speakers represent government, industry and the university, respectively. Our appearance signifies the joint interest of all three of these segments of society in this technological domain of energetic devices embodied with direct intelligent control. But beyond that it signifies that the Robotics Institute has many facets -- that there are many ways to view it.

Thus, my fellow speakers, who have preceded me to the platform, have addressed the Robotics Institute from their own special vantage points. I mean to address it from mine, to ask after the relevance of doing robotics in the university. Within the allotted span this morning there is

time to explore, even in brief, only a single aspect. I have chosen the question of the scientific relevance of robotics. Will robotics contribute to fundamental knowledge?

The question of scientific relevance is of basic concern to the university, which is that institution of our society charged with developing and conserving fundamental knowledge. But note well. Scientific relevance is not the only possible viewpoint, even from the exclusively university perspective I am here adopting. An engineering endeavor is its own intellectual justification, a nexus of problems around which a working life and intellectual curiosity can be organized. Furthermore, what is intellectually exciting encompasses infinitely more than science, or even than science and engineering combined. I select the question of scientific relevance, because it is my own personal passion. I select it also because computer science, with which I have long been associated, is itself a curious mixture of science and engineering, so that the issue has long concerned me.

Each Science has its own Style

We tend to build canonical pictures of how science is done. All of us, scientists and non-scientists alike, are raised on the theory-experiment cycle that is called the scientific method: Start with a theory; from this deduce some consequences; test them by running an experiment; from the results revise the theory. And so repeat the cycle.

Of course, the picture has gradually become more complicated and more realistic. We have all absorbed the viewpoint put forth by Thomas Kuhn that science doesn't proceed gradually, but by a series of conceptual revolutions followed by periods of normal science that assimilate and test the new viewpoints. Still, the presupposition is that all science operates in a given way

My first point, then, is to remind you that science doesn't fit any single mold. We have known it all along, of course. Astronomy was always the canonical exception for the theory-experiment cycle, because it couldn't set up its own experiments in the laboratory and had to rely on pure observation of the stars. However, thinking of exceptions doesn't begin to convey the incredible diversity that exists in how science is accomplished. Every field not only has its own style, but many styles occur within each science. There is Archimedes in his bathtub. There is Baron von Humboldt, tramping up the Orinoco River in 1800 to collect such a multitude of new plant specimens that he takes the rest of his life to analyze the collection and publish it. There is Maxwell, drawing on mechanical models of the ether to help him write his famous equations of electricity and magnetism. There is Einstein, finding the tensor calculus at hand to make possible the mathematical formulation of relativity. There is Watson and Crick, hellbent on solving the problem of the structure of DNA like a giant deliberate puzzle. And on and on

Each of these when analyzed provides a different paradigm for how new knowledge emerges. Each is not isolated, but stands in a chain of other methods of different type. Those who came after Humboldt hardly made progress the same way he did.

The Occurrence of Science in the Midst of Engineering

The emphasis I have just placed on diversity is obviously a lead in to asking how science can occur in the Robotics Institute. I take it as given that robotics is an engineering endeavor. The direct goal is to create artifacts to desirable specifications, in view of human needs and concerns. The underlying intellectual tasks contain large components of design and synthesis (the hallmark of engineering), not just analysis (the hallmark of science).

There are several ways to relate engineering and technology to the development of science. One is that science requires tools, tools of measurement and tools of analysis. It is easy to point to cases where the development of a new instrument or technique has been the direct precursor to scientific advance -- from telescopes, to chromatography, to deep sea drilling gear, to computers. But this is a supporting role. No one doubts that the development of robotic devices will aid the scientific enterprise in a multitude of ways, just as all technological advance has. However, let me set this possibility to one

side in favor of the direct question. Will basic scientific knowledge come out of the engineering study of robotics per se?

The thesis of my talk this morning is that it will. We can expect fundamental science from the Robotics Institute.

This assertion flies in the face of what is perhaps our most fundamental shared model of a uniform method for science. Science is obtained by doing it. That is, to achieve X, do X directly. If you want to know the fundamental nature of matter, or of life, ask yourself directly how to find that out and devote yourself to the task. This view is much more fundamental than science, deriving from our very notion of rational behavior. For instance, it is encased in the old management maxim: If you really want something accomplished, hire a man and make him responsible for doing that and only that.

However, the history of science records many instances when science arose as a side effect of making technological advances. The early development of the steam engine provides a fine example. In 1824, Nicolas Carnot formulated his cycle of heat exchange (known by us, naturally, as the Carnot cycle) in the attempt to understand the low efficiencies of the steam engines of the time. This was such a fundamental advance in understanding that many use it to date the founding of thermodynamics as a distinct discipline.

Another example is the almost romantic tale of the American physicist, Benjamin Thompson, better known as Count Rumford. In the late 18th century, while in charge of boring cannon for his adopted country of Bavaria, he noticed that the cannon continued to heat up as long as boring continued, despite being cooled by water. The then current theories took heat as a fluid (called caloric), that flowed out of the metal when it was bored. Therefore, it should have become depleted. However, the Count showed that, with continued boring, so much caloric must have flowed out, that the cannon would have been melted in the first place. By discrediting the view of heat as a fluid, this took a major step toward our current understanding that heat arises from the motion of molecules.

These cases illustrate how science can emerge even though science itself is not the main goal of the enterprise. Note that these are not examples of "lucky accidents meeting the prepared mind", an oft-noted type of science by side effect. In both examples, the direct goal was to understand a technological process and the new science served that goal. However, it also penetrated deeply enough to add to our fundamental knowledge. This is science arising from a directed engineering endeavor.

What is the Science Lurking in Robotics?

Both my examples above involved heat. Thus, you might think there is something special about heat. There

is: In the 17th and 18th centuries, heat was at the center of developing technologies about which little of fundamental importance was known. Whereas new science would not be expected to arise today in the basic area of thermodynamics at familiar temperatures, all that is necessary is to move to an engineering endeavor where the fundamentals are still unclear. Believe me, robotics is such an area.

From where I stand, it is easy to see the science lurking in robotics. It lies in the welding of intelligence to energy. That is, it lies in intelligent perception and intelligent control of motion. These are to robotics what heat was to the era of Count Rumford and Nicolas Carnot. If we are ever to have devices that operate under their own control to perform tasks that demand large energetic transformations -- whether these be machining, transporting, exploring, mining, constructing, whatever -- they must sense a changing natural environment and control transformations that operate with their own physics in real time.

Now, the scientific task of understanding intelligence is already well under way in artificial intelligence, a subpart of computer science. If an appropriate direct scientific field already exists, why is robotics needed in addition? Why isn't artificial intelligence the supplier of basic scientific capital to robotics, along with other suppliers, such as computer science, mechanical engineering and electrical engineering?

There is an answer, and it further illuminates my basic thesis. Artificial intelligence currently shares with computer science a special view -- it considers information processing divorced from energetics. This creation of an *interior milieu*, in which only information processes occur, is a powerful abstraction, one which helped computer science to emerge by permitting it to focus on the essential mechanisms. But the costs of the abstraction show nowhere more clearly than in the unexplored central problem of robotics -- controlled perceptually coordinated motion. A robot arm contains many degrees of freedom to its motion. To each degree of freedom corresponds some parallel computing power. All these powers must simultaneously be coupled together computationally to operate in real time in a real and mechanically noisy environment. Processing must be intimately interwoven with physics and specialized to it. We know hardly more about this than the 18th century knew about heat.

This example illuminates my thesis, because, though nothing in principle prevents a science from pursuing an area, sciences in fact chose some areas in preference to others. A strong motive is required for artificial intelligence to take as central the scientific problems of perception and motion, with their peculiar patterns of distributed processing. Without such a motive, it is more comfortable -- indeed, it is even scientifically more profitable -- to stay in the interior milieu.

We are now left with the argument about directness. If the purpose is to find out about intelligent motion, why not investigate it directly? If basic science is the goal, the argument will go, attending to all the apparently additional constraints of limited practical devices is just a diversion. Here, I must rely on my Carnots and Rumfords. By working on real problems, phenomena not of the investigator's choosing rise to claim intellectual attention. When much is still unknown, this can be a good thing, and scientific history shows it to be so on many occasions.

The Necessary Condition

From what I have said so far, you might conclude the following. First, a university is concerned (among other things) with basic scientific knowledge. Second, robotics provides (among other things) some chances for science. Therefore, it is ok for an institute of robotics to be at a university. It is as if a happy coincidence has occurred, nothing more.

But I intend more than that. An additional requirement exists for indirect science to emerge from an engineering-oriented effort. The organizational bonds cannot be drawn too tight. If only what is planned for can happen -- or only what is dictated by the goal of overall success -- then little else will occur. There must be the freedom for new problems and new interpretations to emerge and survive while in a nascent and vulnerable state.

By its inherent structure, the modern university is built precisely to provide this freedom. Allow me to put the point in what seems a negative way. A university is incapable of completely controlling any enterprise in which it engages. This happens, neither because we are incompetent nor because we don't understand about management. It happens primarily because we have the twin major goals of education and research. Forever in tension, these keep any single goal from becoming dominant. In the pushing and hauling between their opposing demands, all sorts of freedom exists for the individual, both student and faculty, to pursue his own intellectual aims.

The PhD thesis provides a central example. Theses must produce interesting research results, but they must also teach the student intellectual independence and substance. They cannot be subordinated to an overriding engineering goal. So engineering and performance goals are bent to the needs of education. But, likewise, in engineering environments designs must be completed, the devices must be constructed, and they must really work. So purely educational goals are bent to the needs of productivity and problem solving. Lots of wiggle room exists with all this flexing. This wiggle room, it should be noted, exists entirely within the universe of intellectual concerns, which is common to both research and education.

(continued page 34)

permitting access to a (logically) centralized file system and to other specialized and experimental computing facilities. In addition to a very large address space and high-resolution color display, each machine will have roughly the computing power of a KI-10. SPICE will support a multi-language environment (with interlanguage communication via interprocess communication). There will be two primary programming environments, one based on the algebraic language ADA and the other on LISP.

Scott Fahlman is supervising the development of a complete LISP environment, first by emulating a virtual machine to speed software development, and later experimenting directly with the available hardware

SPICE involves the entire CMU Computer Science Department, both in its development and in its promise of providing the primary computing facility for the department. Though the technical problems are primarily in the area of software systems, we describe it here because a number of the CMU AI researchers are deeply involved in it and because it will form the basis of the AI computing environment in the middle-range future.

Scientific Relevance of Robotics

(continued from page 26)

Many other things happen in universities that increase this diversity of goals. Visitors are welcomed into the environment, often with their own support, hence independence. They attack the problems they see as important from their own interests, which do not coincide entirely either with the educational or the research interests of their host. The result is an increase in

diversity and an increase in the opportunities for science to emerge.

What I claim here is hardly novel. Yet it is important to restate it on this occasion. What often passes for organizational noise in a university effort, is precisely what permits the emergence of the truly fundamental. The noise cannot be too great, but neither can it be too quiet.

Conclusion

Let me conclude. A birth -- or perhaps we should call this a christening -- is a time for good wishes and good hopes. So everyone gets to have his own. It is my good fortune to be a speaker this morning, and thus to share my hopes with all of you. Of all the hopes I could have for the Robotics Institute, the one that is most precious to me -- hence the one I hereby share -- is that it provide the place where some fundamental advances will occur in our scientific understanding of intelligent action. The whole of my remarks have been addressed to showing you why I think that hope is far from being in vain.

I understand fully why others should have other hopes. Hope that the Robotics Institute adds intellectual excitement to the campus. Hope that it creates a vortex of interdisciplinary activity. Hope that the Institute itself, by its very form and function, helps to pioneer a new era of cooperation between university and industry. Hope that it help in our national attempt to become a more productive society. All these hopes, and others too, I also hold. But I must confess, I like mine best. ■

The next issue of AI Magazine will contain a survey of the research being conducted at the Carnegie-Mellon Robotics Institute.

Artificial Intelligence Proceedings

1980: NCAI

Proceedings of the First Annual National Conference on Artificial Intelligence held at Stanford University, sponsored by the American Association for Artificial Intelligence. (In one volume) Order from:

Garcia Robinson, Inc
301 Menlo Oaks
Menlo Park, California 94025

1979: IJCAI6

Held in Tokyo, Japan (In two volumes) Order from:

IJCAI6
Computer Science Department
Stanford University
Stanford, California 94305

1977: IJCAI5

Held at MIT, Cambridge, Massachusetts (In two volumes) Order from:

IJCAI5
Department of Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

1975: IJCAI4

Held in Tbilisi, Georgia, USSR (In two volumes) Order from:

IJCAI4
MIT-AI Laboratory
545 Technology Square
Cambridge, Massachusetts 02139

1973: IJCAI3

Held at Stanford University (In one volume) Order from:

Artificial Intelligence Center
J2044
SRI-International
Menlo Park, California 94025