OFFPRINTS FROM: THE STUDY OF INFORMATION:
INTERDISCIPLINARY MESSAGES
EDITED BY FRITZ MACHLUP & UNA MANSFIELD


# INTELLECTUAL ISSUES IN THE HISTORY OF ARTIFICIAL INTELLIGENCE

Allen Newell

Science is the quintessential historical enterprise, though it strives to produce at each moment a science that is ahistorical. With a passion bordering on compulsion, it heeds the admonition that to ignore the past is to be doomed to repeat it. Science has built its main reward system around discovering and inventing, notions that are historical to the core. Thus, writing about science in the historical voice comes naturally to the scientist.

Ultimately, we will get real histories of artificial intelligence (henceforth, AI), written with as much objectivity as the historians of science can muster. That time has certainly not come. We must be content for a while with connections recorded in prefaces, introductions, citations, and acknowledgments—the web that scientists weave in their self-conscious attempt to make their science into a coherent historical edifice. So far, only a few pieces, such as *Machines Who Think,* provide anything beyond that, and they still have no deliberate historiographic pretensions. [McCorduck, 1979.]

This essay contributes some historical notes on AI. I was induced to put them together originally in response to a request by some of our graduate students in computer science for a bit more historical perspective than is usual in their substantive fare. It is to be viewed as grist for the historian's mill but certainly not as serious history itself. The attempt to define and document all of what I put forward is beyond my resources for the moment. This essay's claim to accuracy, such as it is, rests on my having been a

I thank Elaine Kant and Stu Card for comments on an earlier draft and Paul Birkel and Marc Donner for leading me to write the paper. Note: This research was sponsored in part by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory under Contract F33615-78-C-1551. The views and conclusions contained in the paper are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the United States government.

participant or an observer during much of the period. As is well known to historians, the accuracy of the participant-observer is at least tinged with bias, if not steeped in it. The situation is worse than that; I am not just a participant but a partisan in some of the history here, including parts still ongoing. Reader beware.

## HOW IS THE HISTORY OF A SCIENCE TO BE WRITTEN?

Human endeavors are indefinitely complex. Thus, to write history requires adopting some view that provides simplification and homogenization. The standard frame for the history of science is in terms of important scientific events and discoveries, linked to and by scientists who were responsible for them. This assumes that scientific events declare themselves, so to speak. In many respects this works, but it does so best when the present speaks clearly about what concepts have won out in the end, so that we can work backward through the chain of antecedents, adding only a few dead-ending branches to flesh out the story.

With fields in an early state—and AI is certainly one—critical events do not declare themselves so clearly. Additional frameworks are then useful. Obvious ones of general applicability are proposed theories and research methodologies; neither is very satisfactory for AI. The theoretical ideas put forth have, especially when successful, been embedded in computer systems (usually just as programs but sometimes including special hardware). Often, the systems speak louder than the commentary. Indeed, a common complaint of outsiders (and some insiders) is that there is no theory in AI worthy of the name. Whether true or not, such a perception argues against taking theories as the unit in terms of which history is to be written. As for research methodology, AI as a whole is founded on some striking methodological innovations, namely, using programs, program designs, and programming languages as experimental vehicles. However, little additional methodological innovation has occurred within the field since its inception, which makes for lean history.

Similarly, the more sophisticated units of historical analysis, such as the *paradigms* of Kuhn or the *research programmes* of Lakatos, provide too course a grain. [Kuhn, 1962*a*; Lakatos, 1970.] It can be argued that AI has developed and maintained a single paradigm over its short lifetime, or at most two. Similarly, it has contained at most a small handful of research programmes. But units of analysis work best with enough instances for comparative analysis or for patterns to emerge. There are certainly too few paradigms for an internal history of AI. The same is probably still true of research programmes as well, though it would be of interest to attempt such a description of AI.

Useful frameworks for historical analysis can often be based on the organization of subject matter in a field. AI proceeds in large part by tackling one task after another, initially with programs that can accomplish them crudely, followed gradually by successive refinements. Game-playing, theorem-proving, medical diagnosis—each provides a single developmental strand that can be tracked. Thus, a history of AI as a whole could be written in terms of the geography of tasks successfully performed by AI systems. Almost orthogonal to this task-dimension is that of the intellectual functions necessary for an intelligent system—representation, problem-solving methods, recognition, knowledge acquisition, and so forth—what can be termed the physiology of intelligent systems. All these functions are required in any intellectual endeavor of sufficient scope, though they can be realized in vastly different ways (i.e., by different anatomies), and tasks can be found that highlight a single function, especially for purposes of analysis. Thus, a history can also be written that follows the path of increased understanding of each function and how to mechanize it. Both of these structural features of AI, and perhaps especially their matrix, provide potentially fruitful frameworks for a history. Their drawback is just the opposite from the ones mentioned earlier, namely, they lead to histories that are almost entirely internal, shedding little light on connections between AI and neighboring disciplines.

I settle on another choice, which I will call *intellectual issues*. It is a sociological fact of life that community endeavors seem to polarize around issues—fluoridation versus ban fluoridation, liberal versus conservative. Such polarizing issues are not limited to the purely political and social arena but characterize scientific endeavors as well—heliocentrism versus geocentrism, nature versus nurture. Intellectual issues are usually posed as dichotomies, though occasionally three or more positions manage to hold the stage, as in the tussle between capitalism, socialism, and communism. Intellectual issues are to be distinguished from issues in the real world of action. No matter how complex and ramifying the issues of individual freedom and state control that lie behind a fluoridation campaign, the passage or defeat of an ordinance banning fluoridation is a concrete act and is properly dichotomous. But with nature versus nurture, the dichotomy is all in the eye of the beholder, and the real situation is much more complex (as is pointed out ad nauseum). The tendency to polarization arises from the way people prefer to formulate intellectual issues.

Scientifically, intellectual issues have a dubious status at best. This is true even when they do not have all the emotional overtones of the previous examples. Almost always, they are defined only vaguely, and their clarity seldom improves with time and discussion. Thus, they are often an annoyance to scientists just because of their sloganeering character. Some time ago, in a conference commentary entitled You Can't Play Twenty Questions with Nature and Win, I myself complained of the tendency of cognitive psychology to use dichotomies as substitutes for theories (e.g., serial versus parallel processing, single-trial versus continuous learning). [Newell, 1973*b*.]

Intellectual issues surely play a heuristic role in scientific activity. However, I do not know how to characterize it, nor am I aware of any serious attempts to determine it, though some might exist. Of course, large numbers of scientists write about issues in one way or another, and almost all scientists of an era can recognize and comment on the issues of the day. Were this not true, they could hardly be the issues of the particular scientific day. From a historical and social standpoint, of course, intellectual issues have a perfectly objective reality. They are raised by the historical participants themselves, and both the existence of intellectual issues and the activity associated with them can be traced. They enter the historical stream at some point and eventually leave at some other.

Whether intellectual issues make a useful framework for a scientific history seems to me an entirely open question. Such a history does not at all substitute for histories based on events and discoveries, laid down within a framework drawn from the substantive structure of a field. Still, ever since that earlier paper in 1973, I have been fascinated with the role of intellectual issues. Recently, I even tried summarizing a conference entirely in terms of dichotomies. [Newell, 1980*a*.] Withal, I try it here.

## THE INTELLECTUAL ISSUES

I will actually do the following: I will identify, out of my own experience and acquaintance with the field, all of the intellectual issues that I believe have had some prominence at one time or another. Although I will take the field of AI as having its official start in the mid-1950s, the relevant intellectual issues extend back much earlier. We surely need to know what issues were extant at its birth. I will attempt to put a date both on the start of an issue and on its termination. Both dates will be highly approximate, if not downright speculative. However, bounding the issues in time is important; some issues have definitely gone away and some have come and gone more than once, though transformed each time. I will also discuss some of the major features of the scientific scene that are associated with a given issue. I will often talk as if an issue caused this or that. This is in general illegitimate. At best, an issue is a publicly available indicator of a complex of varying beliefs in many scientists that have led to some result. Still, the attribution of causation is too convenient a linguistic practice to forego.

Table 1 lays out the entire list of intellectual issues. In addition to the short title of the issue, expressed as a dichotomy, there is an indication of an important consequence, although this latter statement is necessarily much abbreviated. The issues are ordered vertically by date of birth and within that by what makes historical sense. All those born at the same time are indented together, so time also moves from left to right across the figure; except that all the issues on hand when AI begins in 1955 are blocked together at the top. Issues that show up more than once are multiply repre-

Table 1. The Intellectual Issues of AI

| | |
|---|---|
| 1640–1945 | Mechanism versus teleology: settled with cybernetics |
| 1800–1920 | Natural biology versus vitalism: establishes the body as a machine |
| 1870– | Reason versus emotion and feeling #1: separates machines from men |
| 1870–1910 | Philosophy versus the science of mind: separates psychology from philosophy |
| 1910–1945 | Logic versus psychologic: separates logic from psychology |
| 1940–1970 | Analog versus digital: creates computer science |
| 1955–1965 | Symbols versus numbers: isolates AI within computer science |
| 1955– | Symbolic versus continuous systems: splits AI from cybernetics |
| 1955–1965 | Problem-solving versus recognition #1: splits AI from pattern recognition |
| 1955–1965 | Psychology versus neurophysiology #1: splits AI from cybernetics |
| 1955–1965 | Performance versus learning #1: splits AI from pattern recognition |
| 1955–1965 | Serial versus parallel #1: coordinate with above four issues |
| 1955–1965 | Heuristics versus algorithms: isolates AI within computer science |
| 1955–1985 | Interpretation versus compilation: isolates AI within computer science |
| 1955– | Simulation versus engineering analysis: divides AI |
| 1960– | Replacing versus helping humans: isolates AI |
| 1960– | Epistemology versus heuristics: divides AI (minor); connects with philosophy |
| 1965–1980 | Search versus knowledge: apparent paradigm shift within AI |
| 1965–1975 | Power versus generality: shift of tasks of interest |
| 1965– | Competence versus performance: splits linguistics from AI and psychology |
| 1965–1975 | Memory versus processing: splits cognitive psychology from AI |
| 1965–1975 | Problem-solving versus recognition #2: recognition rejoins AI via robotics |
| 1965–1975 | Syntax versus semantics: splits linguistics from AI |
| 1965– | Theorem-proving versus problem-solving: divides AI |
| 1965– | Engineering versus science: divides computer science, including AI |
| 1970–1980 | Language versus tasks: natural language becomes central |
| 1970–1980 | Procedural versus declarative representation #1: shift from theorem-proving |
| 1970–1980 | Frames versus atoms: shift to holistic representations |
| 1970– | Reason versus emotion and feeling #2: splits AI from philosophy of mind |
| 1975– | Toy versus real tasks: shift to applications |
| 1975– | Serial versus parallel #2: distributed AI (Hearsay-like systems) |
| 1975– | Performance versus learning #2: resurgence (production systems) |
| 1975– | Psychology versus neuroscience #2: new link to neuroscience |
| 1980– | Serial versus parallel #3: new attempt at neural systems |
| 1980– | Problem-solving versus recognition #3: return of robotics |
| 1980– | Procedural versus declarative representation #2: PROLOG |

sented in the table, according to the date of rebirth, and labeled #1, #2, and so forth. When the ending date is not shown (as in *Reason versus Emotion and Feeling #1: 1870–* ), then the issue still continues into the present.

The issues are discussed in historical order, that is, according to their order in the table. This has the advantage of putting together all those issues that were animating a given period. It has the disadvantage of mixing up lots of different concepts. However, since one of the outcomes of this exercise is to reveal that many different conceptual issues coexisted at any one time, it seems better to retain the purely historical order.

## Mechanism versus Teleology: 1640–1945

We can start with the issue of whether mechanisms were essentially without purpose. This is of course the Cartesian split between mind and matter, so we can take Descartes as the starting point. It is an issue that can not be defined until the notion of mechanism is established. It is and remains a central issue for AI, for the background of disbelief in AI rests precisely with this issue. Nevertheless, I place the ending of the issue with the emergence of cybernetics in the late 1940s. If a specific event is needed, it is the paper by Rosenblueth, Wiener, and Bigelow, which puts forth the cybernetic thesis that purpose could be formed in machines by feedback. [Rosenblueth, Wiener, and Bigelow, 1943.] The instant rise to prominence of cybernetics occurred because of the universal perception of the importance of this thesis. (However, the later demise of cybernetics in the United States had nothing whatsoever to do with any change of opinion on this issue.) AI has added the weight of numbers and variety to the evidence, but it has not provided any qualitatively different argument. In fact, from the beginning, the issue has never been unsettled within AI as a field. This is why I characterize the issue as vanishing with cybernetics. It does remain a live issue, of course, in the wider intellectual world, both scientific and nonscientific, including many segments of cognitive science. Above all, this issue keeps AI in perpetual confrontation with its environment.

Intelligence presupposes purpose, since the only way to demonstrate intelligence is by accomplishing tasks of increasing difficulty. But the relation is more complex the other way around. While purpose could hardly be detected in a device with no intelligence, that is, with no ability at all to link means to ends, no implication follows about the upper reaches of intelligence. Animals, for instance, are obviously purposive yet exhibit strong limits on their intelligence. Thus, settling the question of artificial purpose does not settle the question of artificial intelligence. The continuation of this basic controversy throughout the entire history of AI over whether intelligence can be exhibited by machines confirms this separation. Yet, historically it is not right to posit a separate issue of mechanism versus intelligence to contrast with mechanism versus teleology. No such distinction ever surfaced. Instead, there is an underlying concern about the aspects of mentality that can be exhibited by machines. This shows itself at each historical moment by denying to machines those mental abilities that seem problematic at the time. Thus, the argument moves from purpose in the 1940s to intelligence in the 1950s. With the initial progress primarily in problem-solving, we occasionally heard in the 1960s statements that machines might solve problems but they could never really learn. Thus, the basic issue simply endures, undergoing continuous transformation.

## Natural Biology versus Vitalism: 1800–1920

A critical issue for AI that had come and gone long before AI really began is the issue of vitalism—do living things constitute a special category of entities in the world, inherently distinct from inanimate physical objects. As long as this issue was unsettled, the question of whether the mind of man was mechanical (i.e., nonspecial) was moot. It is difficult to conceive of concluding that the animate world does not generally obey the laws of the physical world but that the mind is an exception and is entirely mechanical. Thus, only if vitalism has been laid to rest for our bodies can the issue be joined about our minds.

The vitalist controversy has a long and well-chronicled history. Retrospectively, it appears as an inexorable, losing battle to find something special about the living, though the issue was joined again and again. Organic matter was just a different kind of matter from inorganic matter—an issue laid to rest finally with the synthesis of urea, an indisputably organic material, from inorganic components in 1828 by Wohler. Organisms had their own inherent internal heat—an issue laid to rest in the work of Bernard by the mid-1800s. For our purposes, the starting and ending dates of the issue are not critical. Vitalism's last champion may be taken to be the embryologist Hans Driesch at the turn of the century, who proposed that organisms develop only by virtue of nonmaterial vital principles, called *entelechies*. [Driesch, 1914.] Issues almost never die, of course, as the continued existence of the Flat Earth Society should remind us. Nevertheless, no substantial intellectual energy has been focused on vitalism in more than fifty years. That the human body is a physical machine, operating according to understood physical laws and mechanisms, sets the stage for considering the mechanistic nature of thought and intelligence.

## Reason versus Emotion and Feeling #1: 1870–

The basic separation of the heart from the head occurred long ago and is a fundamental part of Christian folk psychology. It is background. What concerns us is the ascription of reason (cold logic) to machines and the belief that a machine could have no heart—no feelings or emotions—to ever conflict with its reason. I do not seem to find any good way to fix the initiation of this issue. The striking characteristic of the golem of Rabbi Loew in 1580 seemed to have been literal-mindedness, not heartlessness. And nineteenth-century artificial humans seemed to combine all the human attributes, as did, for instance, Frankenstein's constructed monster. [Shelley, 1818.] But by the twentieth century, certainly in *R.U.R. (Rossum's Universal Robots),* we clearly have the intelligent robot, who is without soul, hence, without emotions or independently felt wants. [Čapek, 1923.] So I have split the latter two dates and taken 1870 as the start.

The relevance of this for AI is in providing a basis for separating machines from humans that is different from the issue of purpose. Although a birthright issue of AI, it does not play a major role. That the issue is there can be seen clearly enough in the paper on "Hot Cognition" by Abelson, which put forth some proposals on how to move machine intelligence in the direction of having affect. [Abelson, 1963.] The lack of prominence stems in part, no doubt, from the strong engineering-orientation of AI, which emphasizes useful mental functions (e.g., problem-solving and learning). In agreement with this, Abelson is one of the few social psychologists associated with AI, and the paper was given at a psychology conference. Thus, this issue remains in the background, waiting to become prominent at some future time.

## Philosophy versus The Science of Mind: 1870–1910

For science as a whole, the separation from philosophy and the acceptance of empiricism as a fundamental tenet occurred centuries ago. For psychology, this occurred very recently, in the last decades of the nineteenth century. Indeed, psychology celebrates the establishment of the first experimental laboratory (Wundt's in Leipzig) in 1879. It was not an especially difficult passage for psychology, given the rest of science as a model. It can be considered complete by the rise of behaviorism, say, by Watson's classic paper. [Watson, 1913.] Thus, this issue emerged and vanished before AI began. The residue was a continuing tradition in philosophy concerned with mind, which was completely distinct from work in psychology and, even more so, from technology. This issue ensured that when AI did emerge, which happened instantly on computers becoming sufficiently powerful,[1] it would be without more than peripheral involvement of the philosophy of mind.

## Logic versus Psychologic: 1910–1945

We continue to lay out the issues—and their resolutions—that were in effect at the birth of AI. This issue concerns whether symbolic logic was to be taken as revealing how humans think or whether humans use some sort of unique "psychologic." It surely started out with logic identified with

[1] A case can be made that serious AI started as soon as computers attained 4K of random-access primary memory.

thought, as Boole's classic monograph entitled *The Laws of Thought* testifies. [Boole, 1854.] But logic was rapidly transformed from an explication of the possible varieties of thinking to a device for probing the foundations of mathematics. We can take the *Principia Mathematica* of Whitehead and Russell as marking the completion of this transformation. [Whitehead and Russell, 1910–1913.] The effect was to separate logic from psychology (and also from the philosophy of mind, although that is a more complex story).

Modern logic, of course, was integrally involved in the development of the digital computer, and, thus, it enters into the history of AI. But logic did not enter AI at all as the logic of thought; that separation remained. Logic was part of the underlying technology of making mechanisms do things. In fact, it was precisely the split of logic from thought that set logic on the path to becoming a science of meaningless tokens manipulated according to formal rules, which, in turn, permitted the full mechanization of logic.

Thus the issue was really settled by 1910, and the status in the first half of the century was that psychologic was not a significant item on the agenda of any science. This, of course, was due to behaviorism's restriction of psychology's agenda. I have placed a date of 1945 for the ending of this issue; this is really an ending of the phase of separating logic from thought. The nerve-net model of McCulloch and Pitts can be used to mark this, along with the work of Turing on which it depended. [Turing, 1936; McCulloch and Pitts, 1943.] They attempted to show that physical systems that echo the structure of the brain could perform all computations, which is to say, all logical functions. Whether this is seen as saying more about the brain or more about logic can be argued; in either case, it brought them back into intimate contact. We might think that the ending of one phase of the issue (the stable separation of logic from thought) should initiate a new phase, namely, a new controversy over the exact nature of the connection. But it did not happen that way. Rather, the issue was not discussed, and basic questions about the mechanization of mind took the form of other issues. The reason that happened cannot be explored here. In part, it comes from the shift with AI from the characterization of the brain in computational terms to the digital computer, where logic played a completely technical and engineering role in describing sequential and combinational logic circuits.

## Analog versus Digital: 1940–1970

When computers were first developed in the 1940s, they were divided into two large families. Analog computers represented quantities by continuous physical variables, such as current or voltage; they were fast, operated simultaneously, and had inherently limited accuracy. Digital computers represented quantities by discrete states; they were slow, operated serially, and had inherently unlimited accuracy. There was a certain amount of skirmishing about which type of computer was better for which type of job. But the technical opinion-leaders maintained a view of parity between the two families—each for its own proper niche. Inevitably, there arose hybrid computers, which claimed to have the best of both worlds: digital control and memory coupled with analog speed and convenience.

It was all over by 1970. The field of computers came to mean exclusively digital computers. Analog systems faded to become a small subpart of electrical engineering. The finish was spelled not just by the increased speed and cost-efficiency of digital systems, but by the discovery of the Fast Fourier Transform, which created the field of digital signal processing and thus penetrated the major bastion of analog computation. The transformation of the field is so complete that many young computer scientists hardly know what analog computers are.

The main significance of this issue, with its resolution, was to help create the discipline of computer science and separate it from electrical engineering. Its effect on AI lies mostly in the loss of an analytical point of view, in which the contrast between analog and digital computation is taken as the starting point for asking what sort of information-processing the nervous system does. An admirable example of this point of view can be seen in the notes for von Neumann's Silliman Lectures, published posthumously. [von Neumann, 1958.] This style of analysis belongs to the world of cybernetics and not to that of AI. I doubt if many young AI scientists have read von Neumann's little book, though it was highly regarded at the time, and von Neumann was one of the towering intellects of the computer field.

## Symbols versus Numbers: 1955–1965

We now come to the first of the issues that characterizes AI itself, as opposed to the background against which it emerged. The digital-computer field defined computers as machines that manipulated numbers. The great thing was, its adherents said, that everything could be encoded into numbers, even instructions. In contrast, scientists in AI saw computers as machines that manipulated symbols. The great thing was, they said, that everything could be encoded into symbols, even numbers. The standard measure of a computation at the time was the number of multiplications it required. Researchers in AI were proud of the fact that there were no multiplications at all in their programs, though these programs were complex enough to prove theorems or play games. The issue was actively pursued as a struggle over how the computer was to be viewed. However, it was joined in an asymmetric way. The bulk of the computer field, and all its responsible opinion-leaders, simply adopted the view that computers are number manipulators. There was no attempt to argue against the view that computers are symbol manipulators. It was just ignored, and the standard interpretation maintained. Researchers in AI, on the other hand, were actively engaged in promoting the new view, considering the standard one to be a radical misreading of the nature of the computer and one that provided a significant barrier to the view that computers could be intelligent.

The result of this clash of views was to isolate AI within computer science. AI remained a part of computer science, but one with a special point of view that made it somewhat suspect, indeed somewhat radical. This isolation is important historically, for it has affected the professional and disciplinary organization of the two fields. It derives ultimately, no doubt, from a basic divergence of views about whether computers can or cannot exhibit intelligence. This overarching issue, of course, continued to be important on its own, as witnessed by the debates that occurred throughout the 1950s on whether machines could think. But the more specific issues that it spawned also had independent lives.

The issue of symbols versus numbers did not arise until after the first AI programs came into existence, circa 1955. Before that time, programs were classified as numerical versus nonnumerical. This latter class was a miscellany of all the things that processed data types other than numbers—expressions, images, text, and so forth.[2] This included the few game-playing and logic programs but much else as well. The symbols-versus-numbers issue emerged only when a positive alternative became formulated, that is, symbolic manipulation. This was not a synonym for nonnumerical processing, for it laid the groundwork for the separation of image- and text-processing from AI. Indeed, the work on machine translation, which started in the early 1950s, was initially considered as one strand in the development of intelligence on machines. [Locke and Booth, 1957.] But that effort became concerned with text and not symbols and developed its own identity as computational linguistics. (All of this, of course, was before text processing in its current meaning emerged—an event that bore no significant relation to the development of computational linguistics.)

I have placed the ending of this issue at about 1965, although I do not have a significant marker event for its demise. The issue is certainly not alive now and has not been for a long time. In part, this is due to the prominence of many nonnumerical data types in computer science generally, such as text and graphics. These make the characterization of computers as number manipulators no longer ring true. In part, it is due to the shift within theoretical computer science to algebraic and logical formalisms, with the concurrent retreat of numerical analysis from its early dominant role. In part, of course, it is due to the success of AI itself and the demonstrations it brought forward of the symbolic character of computation. It is tempting to say that the cause was simply the growth of scientific understanding—but such reasons do not fare well in historical accounts. In any event, my recollection is that the symbols/numbers issue was no longer prominent by the late 1960s, though a little historical digging might place it five years later.

## Symbolic versus Continuous Systems: 1955–

An important characterization of a science, or an approach within a science, is the class of systems it uses to construct its theories. Classical physics, for

[2]The concept of data type did not arrive in clear form until much later.

instance, viewed systems as being described by systems of differential equations. Given a new phenomenon to be explained, a physicist automatically, without a thought, used differential equations to construct his or her theory of that phenomenon. Mathematical psychology in the 1950s and 1960s could be characterized by its acceptance of Markov processes as the class of systems within which to seek theories of particular phenomena.

The issue is within what class of systems should a description of intelligent systems be sought. On one side were those who, following the lead of physical science and engineering, adopted sets of continuous variables as the underlying state descriptions. They adopted a range of devices for expressing the laws—differential equations, excitatory and inhibitory networks, statistical and probabilistic systems. Although there were important differences between these types of laws, they all shared the use of continuous variables. The other side adopted the programming system itself as the way to describe intelligent systems. This has come to be better described as the class of symbolic systems, that is, systems whose state is characterized by a set of symbols and their associated data structures. But initially, it was simply the acceptance of programs per se as the theoretical medium.

Adopting a class of systems has a profound influence on the course of a science. Alternative theories that are expressed within the same class are comparable in many ways, but theories expressed in different classes of systems are almost totally incomparable. Even more, the scientist's intuitions are tied strongly to the class of systems he or she adopts—what is important, what problems can be solved, what possibilities exist for theoretical extension, and so forth. Thus, the major historical effect of this issue in the 1960s was the rather complete separation of those who thought in terms of continuous systems from those who thought in terms of programming systems. The former were the cyberneticians and engineers concerned with pattern recognition; the latter became the AI community. The separation has been strongly institutionalized. The continuous-system folk ended up in electrical-engineering departments; the AI folk ended up in computer-science departments. (It must be remembered that initially computer-science departments were almost exclusively focused on software systems and almost all concern with hardware systems was in electrical-engineering departments.)

I believe this issue largely explains one peculiar aspect of the organization of the science devoted to understanding intelligence: By almost any account, pattern recognition and AI should be a single field, whereas they are almost entirely distinct. By now, in fact, due to another important historical twist, many people in computer science work in pattern recognition. But if such people also know traditional pattern recognition, they are seen as interdisciplinary.

Another interesting implication is buried here. The issue is not properly dichotomous, for there exist other classes of systems within which to search for intelligent systems. One obvious candidate is logic.[3] Were there not

[3] In fact, there are additional possibilities. [Newell, 1970.]

scientists who believed that logic was the appropriate class of systems? And if not, why not? First, by logical systems is meant the class of systems that do logical operations, such as AND, OR, NOT, and so forth.[4] This is the class corresponding to the logic level in the hierarchy of computer structures. The logic level is located between the circuit level and the program (symbol) level. All three levels are equally comprehensive and provide three possibilities for ways of describing intelligent systems. Indeed, circuit and program levels correspond exactly to the continuous and symbol positions of the issue under discussion. Now, in fact, in the early days, there were attempts to build logic machines and discuss the behavior of systems directly in terms of logic circuits. The classical neural networks of McCulloch and Pitts were an effort at modeling the neural system at the logic level. [McCulloch and Pitts, 1943.] But all these efforts rapidly died out and were all but gone by the mid-1960s. My own guess about why this happened is that the hierarchy of computer levels indicated quite clearly what to do with a logic level—namely, compose a higher level system. But this implied simply reproducing existing program-level systems, at least without some new organizational ideas at the program level. But the logic level provided no such ideas, nor could it. Thus, there was nowhere to go. In fact, the history of these efforts seems quite obscure, and tracing the demise of logic as a system language for intelligent systems would be a substantial, though rewarding, undertaking.

## Problem-Solving versus Recognition #1: 1955–1965

An interesting issue grew up in association with the continuous/symbolic split. Those thinking within the framework of continuous systems concentrated on pattern recognition as the key type of task for machines to do—character recognition, speech recognition, and visual-pattern recognition. They also often concentrated on learning (as noted in the following paragraphs), but it was almost always a recognition capability that was being learned. The Perceptron of Rosenblatt can be taken as paradigmatic here. [Rosenblatt, 1958.] Contrariwise, those thinking within the framework of symbolic systems concentrated on problem-solving as the key type of task for machines to do—game-playing, theorem-proving, and puzzle-solving.

This separation of tasks reinforced the split between these groups. To the AI community, the intellectual depth of the tasks performed by the pattern-recognition systems seemed relatively trivial compared with the problem-solving tasks done by the programming systems. But just because of that, a myth grew up that it was relatively easy to automate man's higher reasoning functions but very difficult to automate those functions man shared with the rest of the animal kingdom and performed well automatically, for example,

[4] It might also mean the class of theorem-proving systems using logical calculi; but this is really a subclass of symbol systems.

recognition. Thus, work on recognition was at the foundation of the problem of intelligence, whereas work on problem-solving was an add-on.

The symbolic/continuous split and the problem-solving/recognition split are organically related. Each task is the one most easily approached in terms of the class of systems adopted. However, that does not make the two intellectual issues the same. Scientists can hold quite different attitudes about the two splits, and the two issues can become uncoupled in a different era under different conditions. Both these issues emerged in the late 1950s concurrently with the birth of AI. By 1965 the two fields of AI and pattern recognition had separated rather completely and taken up distinct, relatively permanent institutional roles. The conflict could be considered to have reached a resolution. However, it was to become unstuck again almost immediately.

## Psychology versus Neurophysiology #1: 1955–1965

Strongly coordinated with the issues of symbolic versus continuous systems and problem-solving versus recognition was another, conceptually distinct issue, namely, whether AI would look to psychology or to neurophysiology for inspiration. That human intelligence was to be both guide and goad to engineering intelligent systems was clear. However, this did not discriminate between psychology and neurophysiology. As is well known, these two disciplines speak with entirely separate, though not necessarily contradictory, voices. In general, those concerned with continuous systems and pattern recognition looked to neurophysiology; those concerned with symbolic systems and problem-solving (i.e., AI) looked to psychology. Evidence of the exclusive attention of early AI to psychology (in contradistinction to biology) is amply provided by the two major sets of readings of those years. [Feigenbaum and Feldman, 1963; Minsky, 1968.] By 1965, this issue was no longer a live one, and the cast for AI was set.

The split between neurophysiology and psychology did not dictate the split between symbolic and continuous systems; if anything, it was the other way around. Neurophysiology, of course, was linked to continuous variables, with its signals, networks, and geometry. But experimental psychology was not linked at all to symbolic systems. The dominant class of systems in psychology at the time was that of stimulus/response (S/R) systems, an abstract form of inhibition-and-excitation network. The only alternatives were the continuous fields of Gestalt theory or the pseudo-hydraulic systems of Freudian psychology (both only vaguely defined, though that is irrelevant here). In fact, the class of symbolic systems was discovered within AI and imported into psychology. [Newell and Simon, 1976*a*; Newell, 1980*b*.] Thus, the choice of psychology by AI was made because the class of systems that AI took to work with, that is, programming systems, led to psychologically, not physiologically, revealing tasks.

Neurophysiology played a key role in keeping continuous systems from suffering the same fate as logic systems. Whereas with logic systems there was nowhere to go except toward program-like organizations, with continuous systems there was the brain to model. We need not demand an answer to what the higher organization would be, we could just take as guide the brain as revealed in current neurophysiological work. It is true, of course, that in the late 1940s and early 1950s, the discrete approximation to the nervous system (neurons as digital threshold devices) promised to provide neurophysiological inspiration for the class of logic systems. But under a barrage of criticism, even the engineers came to accept the nervous system as too complex to be modeled by logic-level systems, which is to say, its continuities had to be taken seriously. Thus, without any source of inspiration, logic-level systems faded away as a separate language for modeling intelligence, but continuous systems remained.

## Performance versus Learning #1: 1955–1965

Yet another issue can be identified that is coordinated with the issue of symbolic versus continuous systems. AI concentrated on creating performance systems, that is, systems that performed some task demanding intelligence. Cybernetics and pattern-recognition research concentrated on creating systems that learned. Indeed, another subfield grew up that called itself self-organizing systems. [Yovits, Jacobi, and Goldstein, 1962.] In practice, self-organizing systems largely overlapped with the work in pattern recognition and it had common roots in cybernetics. But self-organizing systems took the problem of learning as the central focus rather than the problem of recognition. For instance, within self-organizing systems, there was considerable interest in embryology, even though it had little to do with recognition at the time.

Through the early 1960s, all the researchers concerned with mechanistic approaches to mental functions knew about each other's work and attended the same conferences. It was one big, somewhat chaotic, scientific happening. The four issues I have identified—continuous versus symbolic systems, problem-solving versus recognition, psychology versus neurophysiology, and performance versus learning—provided a large space within which the total field sorted itself out. Workers of a wide combination of persuasions on these issues could be identified. Until the mid-1950s, the central focus had been dominated by cybernetics, which had a position on two of the issues—using continuous systems and orientation toward neurophysiology—but no strong position on the other two. For instance, cybernetics did not concern itself with problem-solving at all. The emergence of programs as a medium of exploration activated all four of these issues, which then gradually led to the emergence of a single composite issue defined by a coordination of all four dimensions. This process was essentially complete by 1965, although I do not have any marker event. Certainly by 1971, at the second International Joint Conference on Artificial Intelligence in London, it was decided that henceforth the conference would not accept pure pattern-recognition papers, an act which already reflected an existing state of affairs.

### Serial versus Parallel #1: 1955–1965

It is worth noting for future reference that most pattern-recognition and self-organizing systems were highly parallel network structures. Many, but not all, were modeled after neurophysiological structures. Most symbolic-performance systems were serial programs. Thus, the contrast between serial and parallel (especially highly parallel) systems was explicit during the first decade of AI. The contrast was coordinated with the other four issues I have just discussed. However, I do not recollect it playing nearly as active a role as any of the other four, so I have simply added it on as a comment.

### Heuristics versus Algorithms: 1955–1965

These issues were not the only ones that emerged in the first decade of AI's existence, nor the most important. A candidate for the most important initial issue was AI's development of heuristic programs in contradistinction to algorithms. Algorithms were taken to be programs that guaranteed that they would solve a problem or solve it within given time bounds. Good programs were algorithmic, and if not, the fault lay with the programmer, who had failed to analyze his or her problem sufficiently—to know what the program should do to solve this problem. Heuristic programs, on the other hand, were programs that operated by means of heuristic rules of thumb—approximate, partial knowledge that might aid in the discovery of the solution but could not guarantee to do so. The distinction implied that intelligent problem-solving could be attained by heuristic programs. For a short while, one name for the field of AI was heuristic programming, reflecting, in part, a coordination with such subfields as linear programming and dynamic programming (which were also just then emerging).

An important effect of this issue was to isolate AI within computer science but along a different dimension than the issue of symbols versus numbers. Heuristic programming indicates a commitment to a different course than finding the best engineering solution or mathematical analysis of a problem. According to the standard engineering ethos, the proper use of the computer requires the engineer or analyst to exert his or her best intellectual efforts studying the problem, find the best solution possible, and then program that solution. Providing a program with some half-baked, unanalyzed rules seemed odd at best and irrational, or even frivolous, at worst. A good example of this tension can be found in the work of Wang, whose theorem-proving program performed much better than the LOGIC THEORIST. [Newell, Shaw, and Simon, 1957; Wang, 1960.] The thrust of Wang's position was that much better theorem-provers could be built if appropriate results in mathematical logic were exploited. The defense by the AI community stressed finding how humans would solve such problems, in effect denying that the fullest analysis of experimental tasks was the object of the investigation. Another important example was the MACSYMA project to construct an effective computer system for physicists and engineers to do symbolic manipulation of mathematical expressions. Although this work grew out of two prior efforts in AI, it was cast by its leaders as "not part of AI," but, rather, as part of an area of computer science called symbolic manipulation, which took a thoroughgoing engineering and analytical attitude. [Slagle, 1963; Moses, 1967.]

I have put the demise of the issue at the mid-1960s; the issue gradually ceased to be discussed, though the distinction continues to be made in textbooks and introductory treatments. Once the field was underway, with lots of AI systems to provide examples, the point at issue became transparent. Moreover, the distinction has difficulty in being transformed into a technical one, because it is tied to features external to the procedure itself, namely, to the problem that is supposed to be solved and the state of knowledge of the user of the procedure.

## Interpretation versus Compilation: 1955–1985

A third issue served to separate AI from the rest of computer science, in addition to the issues of symbols versus numbers and heuristics versus algorithms. AI programs were developed in list-processing languages, which were interpretive, whereas the mainstream of language development was moving irrevocably toward the use of compilers. Prior to the mid-1950s, programming languages beyond assemblers were interpretive. The major turning point in compilers, FORTRAN, was developed in the mid-1950s,[5] and it determined the direction of programming-language development (though, of course, not without some controversy). Speed of execution was the consideration uppermost in the minds of the programming fraternity. In contrast, AI took the interpretive character of its languages seriously and declared them to be necessary for attaining intelligent systems. This was epitomized by the use of full recursion, but it penetrated throughout the entire philosophy of language design, with the attractive idea of putting intelligence into the interpreter.

This separation of AI programming from mainline high-level language programming, which started immediately at the birth of AI, has persisted to the present. Its effects go much deeper than might be imagined. This separation has played a major role in determining the heavy AI involvement in interactive programming, which contrasts with the minimal involvement of the central programming-languages, with their adherence to the compile-

[5] In fact, the first report of FORTRAN at a scientific meeting occurred at the same session as the first report of a list-processing language. [Backus et al., 1957; Newell and Shaw, 1957.]

and-run operating philosophy. Just for fun, I have indicated the end of this issue in 1985, on the assumption that the coming generation of powerful personal computers will finally force all languages to come to terms with full dynamic capabilities in order to permit interactive programming. But this is pure conjecture, and the separation may now be wide enough to require a generation to heal.

The grounds for this issue can be traced to demands for efficiency on the one hand versus demands for flexibility on the other; perhaps the issue should have been so labeled. For instance, the main programming community in the late 1950s also had a strong negative reaction to list-processing, because of its giving up half the memory just to link the actual data together. But, although the general efficiency issue was always on the surface of discussions, the total situation seems better described in terms of distinct structural alternatives, that is, interpreters versus compilers, list structures versus arrays, and recursion versus iteration.

## Simulation versus Engineering Analysis: 1955–

One issue that surfaced right from the start of AI was whether to make machines be intelligent by simulating human intelligence or by relying on engineering analysis of the task. Those who were primarily trying to understand human intelligence inclined naturally to the simulation view; those who were primarily engineers inclined to the pure task-analysis view. The principle was frequently invoked that we do not build a flying machine by simulating bird flight. On the simulation side, there was more than one position. The majority took the view that casual observation and casual introspection was the appropriate approach—that is, the human was a source of good ideas, not of detail. A few, usually with strong psychological interests or affiliations, took the view that actual experimental data on humans should be examined.

This issue seems never to have produced any important crises or changes of direction in the field; however, it has probably decreased the amount of mutual understanding. There seems to be little movement in a scientist's position on this issue. Each investigator finds his or her niche and stays there, understanding only superficially how those with different approaches operate. The position adopted probably reflects fairly deep attitudes, such as determine whether a scientist goes into an engineering discipline or a social/behavioral discipline in the first place. This is to be contrasted with many fields where methods are effectively neutral means to ends, to be used by all scientists as the science demands. There is little indication of diminution of this issue over the years, although starting in the 1970s, there has been some increase in the general use of protocols to aid the design of AI systems, even when there is no psychological interest.

This completes the set of new issues that arose coincident with the birth of AI. Five of them—symbolic versus continuous systems, problem-solving versus recognition, psychology versus neurophysiology, performance versus learning, and serial versus parallel—separated AI from other endeavors to mechanize intelligence. But the goal of mechanizing intelligence bound all of these enterprises together and distinguished them from the greater part of computer science, whose goal was performing tasks in the service of mankind. Three issues—symbols versus numbers, heuristics versus algorithms, and interpreters versus compilers—clustered together to make AI into a relatively isolated and idiosyncratic part of computer science. Finally one—simulation versus engineering—was purely internal to AI itself.

## Replacing versus Helping Humans: 1960–

An issue that surfaced about five years after the beginning of AI was whether the proper objective was to construct systems that replace humans entirely or to augment the human use of computers. The fundamentally ethical dimension of this issue is evident. Yet, it was not overtly presented as an issue of social ethics but, rather, as a matter of individual preference. An investigator would simply go on record one way or another, in the prefaces of his or her papers, so to speak. Yet, there was often an overtone, if not of ethical superiority, of concordance with the highest ideals in the field. Those whose inclinations were toward AI did not so much meet this issue head on as ignore it. Indeed, it was perfectly possible to take the view that work in AI constituted the necessary exploration for man/computer symbiosis. [Licklider, 1960.]

A relatively weak issue such as this could not really become established unless man/machine cooperation offered technical possibilities and challenges as exciting as constructing intelligent machines. Thus, the beginning of this issue coincides with the appearance of interesting interactive systems, such as SKETCHPAD, which had an immense influence on the field. [Sutherland, 1963.]

Artificial intelligence scientists have had a relatively large involvement in the development of user/computer interaction throughout the history of computer science; for example, in time-sharing in the 1960s and 1970s, in making languages interactive in the 1970s, and in developing personal machines in the early 1980s. One explicit justification given for this involvement was that AI itself needed much better programming tools to create intelligent programs—a reason quite independent of the issue presented here. However, it is not possible to untangle the relations between them without some rather careful historical analysis.

Many of those who opted for working in user/computer cooperation tended not to become part of AI as the latter gradually evolved into a field. However, as I have already noted, it was entirely possible to work in both AI and user/computer cooperation. Still, the net result was an additional factor of separation between those in AI and those in neighboring parts of computer science.

## Epistemology versus Heuristics: 1960–

It is easy to distinguish the knowledge that an intelligent agent has from the procedures that might be necessary to put that knowledge to work to exhibit the intelligence in action.[6] The initial period in AI was devoted almost exclusively to bringing into existence modes of heuristic processing worthy of consideration. In 1959, John McCarthy initiated a research position that distinguished such study sharply from the study of appropriate logical formalisms to represent the full range of knowledge necessary for intelligent behavior. [McCarthy, 1959.] This study was clearly that of epistemology—the study of the nature of knowledge. It bore kinship with the subfield of philosophy by the same name, although, as with so many other potential connections of AI and philosophy, the orientation of the two fields is highly divergent, although the domain of interest is nominally the same.

There has been little controversy over this issue, although the two poles led to radically different distributions of research effort. Work on epistemology within AI has remained extremely limited throughout, although recently there has been a substantial increase. [D. G. Bobrow, 1980.]

## Search versus Knowledge: 1965–1980

In the first years of AI, through the early 1960s, AI programs were characterized simply as highly complex programs, without any particular notion of common structure. For instance, the field was also called *complex information processing* as well as *heuristic programming*. By 1965, however, it had become clear that the main AI programs used the same fundamental technique, which became known as *heuristic search*. [Newell and Ernst, 1965.] This involves the formulation of the problem to be solved as combinatorial search, with the heuristics cast in specific roles to guide the search, such as the selection of which step to take next, evaluation of a new state in the space, comparison of the present state to the posited goal-state, and so on. As the scope of AI programs seemed to narrow, there arose a belief in some AI scientists that the essence of intelligence lay not in search, but in large amounts of highly specific knowledge, or *expertise*. This issue was well enough established by the mid-1970s to occasion the declaration that a paradigm shift in AI had already occurred, the original paradigm having been heuristic search with little knowledge of the task domain and the new paradigm being knowledge-intensive programs. [Goldstein and Papert, 1977.]

It may be doubted that these changes amounted to an actual paradigm *shift*. What clearly did happen was a major expansion of AI research to

[6] Said this way, the connection of this issue to the competence/performance issue discussed later would seem to be overwhelming. However, the research programmes associated with the two issues have never made common cause.

explore systems that included substantial domain-specific knowledge. The subfield currently called expert systems, which includes many of the attempts at constructing applied AI systems, emerged in the mid-1970s in part as a result of this emphasis. However, it became clear that heuristic search invariably continued to show up in these programs. Whenever it did not, the problems being solved by the AI system were extremely easy relative to the knowledge put into the system.

It is useful to see that two types of searches are involved in intelligence. The first is the search of the problem space, that is, heuristic search, which is combinatorial. The second is the search of the system's memory for knowledge to be used to guide the heuristic search. This memory search is through a pre-existing structure that has been constructed especially for the purpose of being searched rapidly; it need not be combinatorial. Both types of searches are required of an intelligent system, and the issue of search versus knowledge helped to move the field to a full consideration of both types. The net result was not so much a shift in the paradigm as a broadening of the whole field. This had become clear enough to the field so that by 1980 the issue can be declared moot.

## Power versus Generality: 1965–1975

Another way to characterize the major early AI programs is that they took a single well-defined difficult task requiring intelligence and demonstrated that a machine could perform it. Theorem-proving, chess and checkers playing, symbolic integration, IQ-analogy tasks, and such management-science tasks as assembly-line balancing—all these fit this description. Again, there was a reaction to this. Although AI could do these sorts of tasks, it could not do the wide range of presumably trivial tasks we refer to as having common sense. The need was for generality in AI programs, not power.

This call had been issued early enough. [McCarthy, 1959.] However, it was really not until the mid-1960s that a significant shift occurred in the field toward the generality and commonsense side. This gave rise to using small constructed puzzles and artificial problems to illustrate various components of everyday reasoning. A typical example was the monkey-and-bananas task, patterned after simple tasks solved by Köhler's chimpanzee, Sultan. Whereas such problems would have seemed insignificant in the early years, they now became useful, because the goal of research was no longer power, but understanding how commonsense reasoning could occur.

By 1975, this shift had run its course, and new concerns for working with relatively large-scale real problems took over with the development of expert systems already mentioned. As could have been expected, the end of this period of emphasis did not mean a shift back to the original issue. Although expert systems tackled real problems and, hence, were obviously powerful, they did not achieve their power by the heuristic-search techniques of the early years; instead they used large amounts of domain-specific knowledge (coupled, sometimes, with modest search).

However, as is usual in the history of science, work on powerful AI programs never stopped; it only diminished and moved out of the limelight. By 1975, highly successful chess programs emerged, built on heuristic-search principles, with an emphasis on large amounts of search—a million positions per move in tournament play—and good engineering. Thus, intellectual issues shift the balance of what gets worked on but rarely shut off alternative emphases entirely.

## Competence versus Performance: 1965–

The Chomskian revolution in linguistics also started in the late 1950s. It was, along with AI, just one of many similar and interrelated developments in engineering, systems, and operational analysis. Although each of these developments had a particularly intense significance for some particular field, for example, linguistics or computer science, they all formed a common interdisciplinary flux. Gradually, these activities sorted themselves into separate subfields or disciplines, developing opposing positions on the issues previously laid out, as we have seen for AI vis-à-vis cybernetics and pattern recognition.

In many ways, linguistics was a special case. It was already a well-formed discipline, and the revolution was at the heart of the discipline, not in some peripheral aspect that could have split off and aligned with other intellectual endeavors. Furthermore, only very few linguists participated in the general flux that was occurring in the world of engineering and applied mathematics. Linguistics was culturally and organizationally quite distinct, having strong roots in the humanities. In fact, it probably made an immense difference that Noam Chomsky became affiliated with the Massachusetts Institute of Technology (MIT).

It was not until the mid-1960s that issues emerged that determined relations between linguistics and other subfields and disciplines. A principal issue was the distinction between competence and performance, which was moved to a central position in the new linguistics by Chomsky. [Chomsky, 1965.] Linguistic competence was the general knowledge a speaker had of the language, in particular, of the generative grammar of the language. Performance was the actual production of utterances, which could be affected by many additional factors, such as cognitive limits, states of stress, or even deliberate modifications for effect. The distinction made useful operational sense for linguistics, because there were two sources of evidence about human-language capabilities, the actual utterance and the judgment of grammaticality—a sort of recall/recognition difference, although that analogy was never exploited.

This distinction might seem innocuous from the standpoint of science history, that is, purely technical. In fact, it served to separate quite radically the sciences concerned primarily with performance, namely AI, computational linguistics, cognitive psychology, and psycholinguistics, from linguistics proper. Linguistics itself declared that it was not interested in performance. More cautiously said, competence issues were to have absolute priority on the research agenda. But the effect was the same: Work in any of the performance fields was basically irrelevant to the development of linguistics. There could be a flow from linguistics to these other fields, and, indeed, there was an immense flow to psycholinguistics, but there could not be any significant flow in the other direction.[7]

A more effective field-splitter would be hard to find. It has remained in effect ever since, with the competence/performance distinction being extended to other domains of mentality. This has certainly not been the only significant cause of the separateness of AI from linguistics. There are important isolating differences in method, style of research, and attitudes toward evidence. Many of these other issues share substance with the competence/performance distinction and affect the separation between psychology and linguistics much more than that between AI and linguistics. Thus, perhaps these issues can be left to one side.

## Memory versus Processing: 1965–1975

During the immediate postwar decades, the mainstream of individual human psychology was strongly influenced by the general ferment of engineering, system, and operational ideas (as I have previously termed it). This involved human factors and information theory in the early 1950s; and signal-detection theory, control theory, game theory, and AI in the mid-1950s. As with linguistics in the period of 1955–1965, all these ideas and fields seemed to mix while matters sorted themselves out. By the mid-1960s, psychology had focused on memory as the central construct in its view of man as an information processor. Short-term memory and the visual iconic store combined to provide an exciting picture of the interior block-diagram of the human mental apparatus (what would now be called the architecture). This settled what the main lines of investigation would be for the field; the marker event for this conviction is Neisser's book, *Cognitive Psychology*. [Neisser, 1967.]

This settlement is important for the history of AI, because AI's influence on psychology in the 1955–1965 period was primarily in the area of problem-solving and concept formation. With psychology opting for memory structure, psychology and AI went fundamentally separate ways. Although the work on problem-solving remained a common concern, it was a sufficiently minor area in psychology, so that it exerted only a modest integrating effect. AI itself during this period had little interest in memory structure at the block-diagram level. Psychologically relevant research on memory by AI researchers did exist but moved out of AI into psychology; for example, the

[7]This is not the whole story of the relations of linguistics with other fields; for example, there have been important contacts with logic and philosophy.

work on EPAM (Elementary Perceiver and Memorizer). [Simon and Feigenbaum, 1964.]

In the second half of the 1960s came another major advance in cognitive psychology, namely, the discoveries of how to infer basic processes from reaction times. [Neisser, 1963; Sternberg, 1966.] This insight promised even greater ability to dissect human cognitive processes and confirmed the basic choice of psychology to analyze the block-diagram level of cognition. This insight also broadened the analysis from just memory structure to the stages of information-processing. In this respect, it might seem better to call the issue under discussion one of system levels: AI focusing on the symbolic level and psychology focusing on the architecture,[8] that is, the equivalent of the register-transfer level. However, the concern with memory so dominates the years prior to 1965, when this issue was being sorted out, that it seems preferable to label it memory versus processing.

Long-term memory has been absent from the previous account. During this period, AI was certainly concerned about the structure of long-term memory, under the rubric of semantic memory. This would seem to provide common ground with psychology, yet initially it did not do so to any great extent. Two factors seem to account for this. First, in psychology, the new results, hence the excitement, all involved short-term memories. The established theory of learning, interference theory, against which these new ideas about memory made headway, assumed a single memory, which was in essence long-term memory. Second, the memory that psychology considered was episodic—learning what happened during an episode, such as learning what familiar items were presented at a trial. This stood in marked contrast with semantic memory, which appeared to be a timeless organization of knowledge. Only gradually did the psychologically relevant work on semantic memory by a few investigators capture any significant attention within cognitive psychology. The seminal publication of Anderson and Bower's *Human Associative Memory* can be taken as a marker of the beginning of this attention. [Anderson and Bower, 1973.]

### Problem-Solving versus Recognition #2: 1965–1975

In 1965, AI took back the problem of recognition that had become the intellectual property of the pattern-recognition community. This can be marked rather precisely by the work of Roberts on the recognition of three-dimensional polyhedra. [Roberts, 1965.] The essential features were two: First, recognition was articulated, that is, the scene had to be decomposed or segmented into subparts, each of which might need to be recognized to be a different thing. Thus, the result of recognition was a description of a scene rather than just an identification of an object. But a description is a symbolic

[8] Although the term *architecture* is just now coming into common use in psychology.

structure that has to be constructed, and such processes were quite outside the scope of the pattern-recognition techniques of the time, though exactly of the sort provided by AI. Second, a major source of knowledge for making such recognitions came from adopting a model of the situation (e.g., it consists only of polyhedra). This made recognition processes strongly inferential, again fitting in well with work in AI, but not with work in pattern recognition.

By the late 1960s, work on vision was going on throughout AI, but the transformation went further than just vision. Three laboratories (at MIT, Stanford, and the Stanford Research Institute) started major efforts in robotics. Vision was to be coupled with arms and motion and in at least one AI center (Stanford), with speech. The entire enterprise was radically different in its focus and problems from the research in pattern recognition that was still going on in parallel in departments and research centers of electrical engineering. In fact, there was little actual controversy to speak of. Both groups simply did their thing. But likewise, there was no substantial rapprochement.

## Syntax versus Semantics: 1965–1975

The Chomskian revolution in linguistics was strongly based on theory. Built around the notions of generative and transformational grammar, it posited three distinct components (or modules) for phonology, syntax, and semantics, each with its own grammar. The initial emphasis was on syntax, with work on semantics much less well developed.[9] Despite cautions from the competence/performance distinction, the inference was clear from both the theory and practice of linguistics—syntactic processing should occur in a separate module independently of semantic processing. Indeed, what computational linguistics there was in association with the new linguistics involved the construction of programs for syntactic parsing.

In the late 1960s, a reaction to linguistics arose from within the AI and computational linguistics communities. It took the form of denying the separation of syntax and semantics in the actual processing of language. The initial analysis of an utterance by the hearer was as much a question of semantics as of syntax. Language required an integrated analysis by the hearer and, hence, by the theorist. This reaction can be marked by the work of Quillian, whose introduction of semantic nets was a device to show how semantic processing could occur directly on the surface structure of the utterance (though presumably in conjunction with syntax). [Quillian, 1968.]

This reaction was grounded more broadly in the assertion of the importance of processing considerations in understanding language, the very thing

[9] There was work on phonology, but the domain lay outside the range of interest of AI and, in fact, of psychology as well.

denied by the competence/performance distinction. It sought to put processing considerations into the mainstream of linguistic studies, the latter being owned, so to speak, by the linguistics community. One result, as might have been expected, was to compound the separation between linguistics, on the one hand, and computational linguistics and AI, on the other. Another was to create a stronger independent stream of work on language in AI with its own basis.

## Theorem-Proving versus Problem-Solving: 1965–

Theorem-proving tasks have always been included in the zoo of tasks studied by AI, although the attention these tasks received initially was sporadic. However, some logicians and mathematicians worked on theorem-proving in logic, not just as another task, but as the fundamental formalism for understanding reasoning and inference. In the last half of the 1960s, with the development of a logical formalism called resolution, this work in theorem-proving took center stage in AI. [Robinson, 1965.] It seemed for a time that theorem-proving engines would sit at the heart of any general AI system. Not only was their power extended rapidly during this period, but a substantial amount of mathematical analysis was carried out on the nature of theorem proving in the predicate calculus. Even further, theorem-proving programs were extended to handle an increasing range of tasks, for example, question-answering, robot-planning, and program-synthesis.

A consequence of this success and viewpoint was that theorem-proving was taken to be a fundamental category of activity distinct from other problem-solving, with its own methods and style of progress. A good indicator of this is Nilsson's AI textbook, which divides all problem-solving methods of AI into three parts: state-space search, problem-reduction (i.e., subgoals), and predicate-calculus theorem-proving. [Nilsson, 1971.] It is not clear whether this issue has been laid to rest by now or not. As recounted in the following section, under the procedural/declarative issue, theorem-proving has become much less central to AI since the mid-1970s. But theorem-proving and problem-solving still remain distinct research strands.

## Engineering versus Science: 1965–

Computer science is torn by a fundamental uncertainty over whether it is an engineering or science discipline. There is no doubt about the engineering side; computer science designs and creates artifacts all the time. The doubt exists on the nature of the science involved. Computer science certainly studies intellectual domains that are not part of other disciplines. The question is whether or not they have the character of a science. However, the dichotomy need not be accepted: A third alternative is that the unique intellectual domain of computer science is part of mathematics. Computer science would then join other engineering specialties, such as control theory and information theory, which have their own characteristic mathematical development.

Much rests on the putative outcome of this issue: What should computer science be like in the future? Should departments of computer science be part of the college of engineering or the college of arts and sciences? What status should be accorded to various subdisciplines in computer science? Can a thesis involve just a design? And more. The start of this issue coincides with the creation of departments of computer science in the mid-1960s, which served to raise all these questions. Whether the issue will ever be laid to rest is unclear, but it is certainly unlikely while the whole field grows dynamically, with a continuing flood of new and destabilizing notions.

Artificial intelligence participates along with the rest of computer science in the uncertainties over whether it is an engineering or science discipline. However, the issue for AI has its own special flavor. AI participates with many disciplines outside computer science in the attempt to understand the nature of mind and intelligent behavior. This is an externally grounded scientific and philosophic goal, which is clearly not engineering. Thus, the nature of the science for AI is not really in doubt as it is for the rest of computer science. However, this does not end the matter, for interactions occur with other issues. For instance, to the extent that we are oriented toward helping humans rather than replacing them, we may not wish to accept the understanding of the nature of mind as a scientific goal, but only as a heuristic device.

The orientation toward engineering or science can have major consequences for how a field devotes its energies. Currently, for example, an important divergence exists in the subfield of computer vision. Should the nature of the environment be studied to discover what can be inferred from the optic array (a scientific activity); or should experimental vision systems be constructed to analyze the data they generate within the framework of the system (an engineering activity)? That both activities are legitimate is not in question; which activity gets the lion's share of attention is in dispute. And there is some indication that an important determiner is the basic engineering/science orientation of a given investigator.

## Language versus Tasks: 1970–1980

The 1970s saw the emergence of concerted efforts within AI to produce programs that understand natural language, amounting to the formation of a subfield, lying partly in AI and partly in computational linguistics. The key markers are the works of Woods and Winograd. [Woods, 1970; T. Winograd, 1971]. This issue had been building for some time, as we saw in the issue of syntax versus semantics.

The emergence of such a subfield is in itself not surprising. Natural language is clearly an important, even uniquely important, mental capability. In addition to AI, there existed another relevant field, computational linguistics, concerned generally with the application of computers to linguistics. Neither is it surprising that this subfield had almost no representation from linguistics, although, of course, linguistics was of obvious central relevance.[10] The syntax/semantics issue, which had reinforced the separation of linguistics from AI, was a primary substantive plank in the programme of the new subfield.

What is interesting was the creation of another attitude within a part of AI, which can be captured by the issue of language versus tasks. Studying the understanding of language was seen as a sufficient context for investigating the nature of common sense. An important discovery was how much knowledge and inference appeared to be required to understand even the simplest sentences or short stories. Thus, the very act of understanding such stories involved commonsense reasoning and, with it, the essence of general human intelligence. Programs could be interesting as AI research, so the attitude went, without doing any other task in addition to understanding the presented language input. The effect of this strategic position was to separate the work in natural-language processing from the tradition in AI of posing tasks for programs to do, where the difficulty could be assessed. The issue did not occasion much discussion, although its effects were real enough. The issue was masked by the fact that understanding by itself was a difficult enough task for AI research to make progress on. No one could object (and no one did) to not adding what seemed like an irrelevant second difficult task for the system, which would simply burden the research endeavor.

### Procedural versus Declarative Representation #1: 1970–1980

Recall that resolution theorem-proving flourished in the late 1960s and bid fair to become the engine at the center of all reasoning. In fact, it took only a few years for the approach to come up against its limitations. Despite increases in power, relative to prior efforts, theorem provers were unable to handle any but trivial tasks. Getting from logic to real mathematics—seen always as a major necessary hurdle—seemed as far away as ever.

The reaction to this state of affairs became known as the procedural/declarative controversy. Theorem provers were organized as a large homogeneous database of declarative statements (clauses in resolution), over which an inference engine worked to produce new true statements to add to the database. This was the essence of a declarative representation of knowledge and its attractions were many. Its difficulty lay in the costs of processing. The inference engine treated all expressions in the database alike or, more precisely, without regard for their semantics. There also seemed no

[10] Among the contributors to the first conference on Theoretical Issues in Natural Language Processing, a series that became the forum for this subfield, I can identify only one mainstream linguist. [Schank and Nash-Webber, 1975.]

way for a theorem prover to be given information about how to solve problems. These two features added up to a major combinatorial explosion. The remedy—the procedural side of the issue—lay (so it was claimed) in encoding information about the task in procedures. Then knowledge would be associated directly with the procedures that were to apply it; indeed, the procedures would embody the knowledge and, thus, not have to be interpreted by another inference engine. This would permit the appropriate guidance for problem-solving and, thus, keep the combinatorial explosion under control.

There are irremediable flaws in both sides of the argument whether knowledge should be coded in procedural or declarative form, just as there are irremediable flaws in both sides of the argument whether a program is heuristic or algorithmic. Both procedural and declarative representations are necessary to make any computation at all happen. In consequence, arguments over the issue were largely inconclusive, although they produced the closest thing to a public issue-controversy in AI's short history. However, the effect on the course of AI research was enormous. First, work on theorem-proving shrank to a trickle, with what remained mostly devoted to nonresolution theorem-proving. Second, so-called planning languages emerged as a result—PLANNER, QA4, CONNIVER, POPLAR, and so forth. [Bobrow and Raphael, 1974.] These programming-language systems were intended to provide a vehicle for writing the sorts of domain-dependent, procedure-oriented theorem provers called for in the debate. While that did not quite happen, these languages in themselves provided a major conceptual advance in the field. The effects of this issue had about run their course by 1980.

## Frames versus Atoms: 1970–1980

In a paper that circulated widely before it was published in the mid-1970s, Marvin Minsky raised the issue about the size of representational units in an intelligent system. [Minsky, 1975.] Knowledge should be represented in *frames,* which are substantial collections of integrated knowledge about the world, rather than in small atoms or fragments. The basic issue is as old as the atomistic associationism of British empiricism and the countering complaints of the Gestaltists. How are the conflicting requirements for units of thought and contextual dependence to be reconciled?

This issue had hardly surfaced at all in the first decade of AI. List structures, the basic representational medium, were in themselves neither atomistic nor wholistic but adaptable to whatever representational constructs the designer had in mind.[11] But the coming to prominence of resolution-theorem-

[11] This is because list structures approximate general symbolic systems. The neutrality is easily confirmed in the continued and universal use of list-processing languages to realize systems of all kinds along this dimension.

proving in the late 1960s brought with it as a side effect the *clause* as the unit of representation. The clause was a primitive assertion that could not be broken down into a conjunction of other assertions—primitive predicates *P*, negations of primitive predicates ~*P*, disjunctions *P* or *Q*, implications *P* implies *Q*, and so forth. The total knowledge of the system was to be represented as the conjunction of clauses—that is, to use the old Gestaltist phrase, as an *And-sum* of separate bits of knowledge.

Thus, the issue of size of representational unit grew out of the same ground as the procedural versus declarative controversy, and, indeed, it was articulated by the same group at MIT who had made most of the latter issue. As is always the case, concern was, in fact, widespread but had been subordinated to other concerns. [Abelson, 1973; Norman, 1973; Schank, 1973.] Minsky was the first one to give clear voice to the concern. The effect of the paper was dramatic, despite the fact that the paper itself was entirely speculative and discursive. Throughout AI, the concept of the frame as the appropriate data structure was widely embraced. By 1980, frame systems were an established part of AI, and a very substantial fraction of the work in knowledge representation was involved in such systems.

Much follows on this development (in conjunction with the procedural/declarative issue)—the rise of substantial research effort in knowledge representation and the strengthening of renewed ties with philosophy. [Brachman and Smith, 1980.] These efforts conjoin with those of AI epistemology, discussed earlier. They raise some new issues, such as the relation of philosophic work on meaning to directly inspired computational models. But these issues have not yet jelled enough to be included in their own right.

## Reason versus Emotion and Feeling #2: 1970–

Philosophy has a long-standing concern with the mechanization of mind. Indeed, under the rubric of the mind/body problem, it can be said almost to own the problem, it having been bequeathed to philosophy by Descartes. In its genesis, AI had very little involvement with philosophy, beyond the background awareness that comes from participation in the general intellectual culture. No philosophers of mind were involved and no technical philosophical issues were dealt with. A glance at the content of the two fields provides one obvious clue. The phenomena attended to in philosophy are sensations as subjective experiences—*raw feels,* to use a bit of philosophic jargon. A typical article is entitled "The Feelings of Robots." [Ziff, 1959.] Thus, though AI and philosophy of mind ostensibly deal with the same problem, in fact they go after largely distinct phenomena.[12]

[12] Another example is the problem of induction, where philosophy is concerned with the certainty of induction and AI is concerned with performing the inductions. [Newell, 1973c.]

The issue has not been especially active, but it has been raised. [Gunderson, 1971.] It is argued that performance functions (i.e., those functions AI currently deals with, called *program-receptive* functions) can be mechanized; but that sentient functions (i.e., feelings, called *program-resistant* functions) cannot. Whether this will ever grow to a substantial controversy is hard to tell at this point. It is certainly available as a reserve position that can serve to separate AI from the philosophy of mind. It adds to the general background concern, discussed in the first occurrence of this issue, of the absence of emotion and feeling in the development of intelligent systems.

## Toy versus Real Tasks: 1975–

As noted in the power/generality issue, the field took a shift in the mid-1960s away from powerful programs toward programs that could exhibit common sense. Further, as noted in the language/tasks issue, this line further transmuted to being concerned with understanding via the understanding of natural language. Concomitantly, programs were often built to work on small simple illustrative tasks or environments, usually puzzles or made-up situations.

By the mid-1970s some systems had been developed that worked with real tasks that had substantial intellectual content, to judge from their role in the real world. The initial such system can be taken to be DENDRAL, which determined the structural formula for chemical molecules, given the data on the mass spectrogram.[13] [Lindsay, Buchanan, Feigenbaum, and Lederberg, 1980.] DENDRAL began in the late 1960s and grew in power throughout the early 1970s. It was joined in the mid-1970s by several systems that performed competently in real medical-diagnosis tasks, of which MYCIN was the paradigm. [Shortliffe, 1974.] This was the immediate locus of expert systems, which, as previously noted, grew up as part of the general emphasis on knowledge in contrast to search. With it grew an attitude that AI in general should no longer work on small illustrative, artificial tasks but that it was time to work on real tasks. The simple artificial tasks came to be called toy tasks, not just because the term conveys the contrast between childish and grown-up pursuits, but also because stacking children's blocks had become a favorite illustrative task environment.

The tension between basic research and application exists in all sciences at all times. Sciences sometimes build institutional structures to contain the tension. As we saw in the issue of science versus engineering, computer science has kept its basic and applied components mixed together in a single discipline, thus exacerbating the tension. The tension was, in fact, especially

[13] The other system often mentioned similarly is MACSYMA, the highly sophisticated program at the Massachusetts Institute of Technology for doing symbolic mathematics. As mentioned earlier, it had deliberately removed itself from being an AI program.

severe for AI during the decade of the 1970s. The climate in Washington was not benign for basic research in general, and there was sustained pressure from AI's primary government funding agency (DARPA—Defense Advanced Research Projects Agency) to make AI pay off. That said, however, the distinction between toy versus real tasks is not solely the distinction between basic and applied research. Tasks taken from the real world and performed by intelligent humans as part of their working lives carry a prima facie guarantee of demanding appropriate intelligent activity by systems that would perform them. It can be argued that such tasks are the appropriate ones for AI to work on, even if the goal is basic research. Thus, the toy-versus-real-tasks issue stands ambiguously for both meanings—basic versus applied and irrelevant versus relevant basic science.

### Serial versus Parallel #2: 1975–

By the mid-1970s, computer science had for some time been seriously exploring multiprogramming and multiprocessing. These provided the groundwork for considering parallel systems for doing AI. A major instigation occurred with the development of the Hearsay-II model of speech understanding. [Lesser and Erman, 1977.] Hearsay-II comprised a number of knowledge sources (acoustic, phonetic, phonological, lexical, syntactic, semantic, and pragmatic), each working concurrently and independently off a common blackboard that contained the current working state about the utterance and each contributing their bit to the evolving recognition and reacting to the bits provided by the others.

The Hearsay-II structure was certainly a parallel one, but it was at a level of parallelism quite different from earlier network models, namely, a modest number (tens) of functionally specialized processes. Furthermore, individual processes remained fundamentally symbolic (even though lots of signal-processing was inherent in the speech-recognition task). Hearsay-II was only one of several efforts to pursue the notion that an intelligent system should be thought of in terms of communicating subprocesses rather than as an individual serial machine. A metaphor arose for thinking about an intelligent system—the scientific community metaphor—which took the operation of science, with its notion of cooperation, publication, experiment, criticism, education, and so forth, as the appropriate model for intelligent activity. Gradually, a group of people emerged interested in working on distributed AI.

### Performance versus Learning #2: 1975–

As noted earlier, learning was generally associated with work on pattern recognition. With the split between problem-solving and recognition, work on learning within AI declined. As always, it never stopped entirely. Indeed, such is the basic fascination with learning processes, and with the belief that they hold the key to intelligence, that each learning program that was constructed received substantial attention.[14] [Samuel, 1959; D. A. Waterman, 1970; Winston, 1970; Sussman, 1975.] However, each learning system was relatively idiosyncratic, with its own interesting lessons, so that the whole did not add up to a coherent effort for the field.

A reversal of this state of affairs developed by the late 1970s. It was triggered by the spread of a class of programming systems, called production, or rule-based systems, which are used for both constructing expert systems and analyzing human cognition. [Waterman and Hayes-Roth, 1978.] To appreciate their role in the resurgence of work on learning, we must take a step back. To create a learning system requires solving two research problems. First, a space of potential performance programs must be created, in which learning will constitute moving from one program to another, searching for programs with better performance. If the space of programs is too vast and irregular, then learning is, in effect, automatic programming, and it becomes extremely difficult. If the space is too limited, then learning is easy, but the performance programs are of little significance. Determining the right space is, thus, a critical research activity. Second, given the space, it is still necessary to design an interesting learning system, for the space only lays out the possibilities. Thus, inventing the learning system is also a critical research activity. A major reason why early AI learning-systems seemed so idiosyncratic was that each made unique choices on both these dimensions. Most important, doing research on learning was doing a double task and taking a double risk.

A production system is composed entirely of a set of *if-then* rules (if such and such conditions hold, then execute such and such actions). At each instant, the rules that hold are recognized, and a single rule is selected to execute. In such a system, the natural space of performance programs consists of subsets of if-then rules, and the primitive act of learning is to add a new rule to the existing set (or sometimes to modify an existing rule in some simple way, such as by adding another condition). This space of performance programs is neither too limited nor too open, since it is easy to restrict the rules to be learned to a special class. As a consequence, the first research choice is essentially made for the researcher, who can then concentrate on constructing an interesting learning program. Moreover, learning programs will have much in common, since they now use similar spaces of performance programs. Indeed, this is just what happened in the late 1970s as researchers began to construct a wide variety of small learning systems, all built around variants of the production-system formalism. [Michalski, Carbonell, and Mitchell, 1983.] It must be realized, of course, that such focusing of effort does not remove the collective risk. If production systems

[14] Some other systems were built, which might have been viewed as learning systems, but, instead, were taken simply to be performance programs in specialized task environments, for example, induction programs.

are the wrong program organization to be exploring, then the entire field is moving down an unproductive path.

### Psychology versus Neuroscience #2: 1975–

AI would appear to be at the mercy of the immense gulf that continues to separate psychology and the biology of the brain. As each field continues to progress—which both do dramatically—hopes continually spring up for new bridging connections. No doubt at some point the permanent bridge will be built. So far, although each increment of progress seems real, the gap remains disappointingly large.

It is possible that AI has a major contribution to make to this by exploring basic computational structures at a level that makes contact with neural systems. In the early instance of psychology versus neurophysiology (which was before the term *neuroscience* had been coined), that possibility seemed quite remote. The theoretical structures that did make contact with neurophysiology were remote from the computational structures that preoccupied AI researchers. Then the split occurred, with pattern recognition all but moving out of computer science.

In the mid-1970s, a new attempt began to connect AI with neuroscience, initiated by the work of David Marr. [Marr, 1976.] The emphasis remained on vision, as it had been in the earlier period. But the new effort was explicitly computational, focusing on algorithms that could perform various low-level vision functions, such as stereopsis. Although Marr's effort was new in many ways, and based on specific technical achievements, most of the global issues of the earlier time reappeared. This work has now expanded to a larger group, which calls its work, among other things, the new connectionism, and promises to be a substantial subfield again, this time within AI.

### Serial versus Parallel #3: 1980–

The new wave of neuroscience-inspired AI contains, of course, a commitment to highly parallel network structures. The issue of serial versus parallel merits a separate entry here to maintain a clear contrast with the distributed AI effort, which defined the second wave of concern with parallel systems. In this third phase, the degree of parallelism is in the millions, and computing elements in the network have modest powers; in particular, they are not computers with their own local symbols. In the new structures, computation must be shared right down to the roots, so to speak. The interaction cannot be limited to communicating results of significant computations. Furthermore, the communication media between elements are continuous signals, and not just bits. However, unlike the earlier work, these new computational systems are not to be viewed as neural nets; that is, the nodes of the network are not to be put in one-to-one correspondence with neurons, but, rather, with physiological subsystems of mostly unspecified character.

### Problem-Solving versus Recognition #3: 1980–

Robotics has returned to AI after having left it for most of the 1970s. Perhaps it is unfortunate to call the issue problem-solving versus recognition, since recognition is only one aspect of robotics. The main sources of the new wave of effort are external to AI—industrial robotics plus the concern of the decline in American productivity and the trade position of the United States vis-à-vis Japan and West Germany. The initial growth of industrial robotics took place largely outside of AI as a strictly engineering endeavor. As a result, the initial growth tended to minimize the intelligence involved, for example, sensory-motor coordination. One component of the new association of robotics with AI is the coupling of significant amounts of vision with manipulators, reflecting the continued advance of vision capabilities in AI throughout the 1970s. (Touch and kinesthetic sensing is increasingly important, too, but this does not build so strongly on prior progress in AI.) Importantly, along with industrially motivated aspects, there is also a revival of basic research in manipulation and movement in space and over real terrains.

It might seem that this is just another purely technical progression. But with it has returned, as night follows day, the question of the relation of AI and robotics as disciplines, just as the question was raised in the issue of problem-solving versus recognition during the late 1960s. Is robotics a central part of AI or only an applied domain? Do graduate students in AI have to understand the underlying science of mechanics and generalized coordinate systems that are inherent in understanding manipulation and motion? Or is that irrelevant to intelligence? Cases can be made either way. [Nilsson, 1982.]

### Procedural versus Declarative Representation #2: 1980–

In the late 1970s, a new programming system called PROLOG emerged, based on resolution-theorem-proving and constituting, in effect, a continuation of the effort to show that declarative formulations can be effective. [Kowalski, 1979.] The effort is based primarily in Europe, and it is a vigorous movement. The attack is not occurring at the level of planning languages, but at the level of LISP itself. Over the years, LISP has established itself as the lingua franca of the AI community. Even though various other programming systems exist, for example, rule-based systems of various flavors, practically everyone builds systems within a LISP programming environment. The planning languages (PLANNER, CONNIVER, etc.), which showed how to effect another level of system organization above LISP, have not proved highly effective as a replacement, and they receive only modest use. As already noted, their contribution has been primarily conceptual. Thus, although the original attack on theorem-proving was in terms of the planner languages, the modern counterattack is at the level of LISP. By being centered in Europe, with very little attention paid currently to

PROLOG in the major AI centers in the United States, the issue takes on additional coordinated dimensions. The outcome is far from clear at this juncture.

## DISCUSSION

It should be clear by now why I entered the caveats about historical accuracy at the beginning. Each of the issues raises serious problems of characterization and historical grounding. No attempt has been made to define an intellectual issue, so that some modestly objective way could be found to generate a complete set of issues, for example, by placing a grid over the literature of the field. Several additional issues might well have emerged, and some of those presented here might not have made the grade. Thus, the population of issues exhibited must be taken, not just with a pinch of salt, but soaked in a barrel of brine. Similar concerns attend dating the issues and my interpretation of them; nevertheless, some comments about the total picture seem worthwhile.

### What Is Missing?

I do know why some issues did not make it. Three examples will illustrate some reasons. The first is the broad but fundamental issue of the ethical use of technology and the dehumanization of people by reduction to mechanism. This issue engages all of technology and science. It seems particularly acute for AI, perhaps, because the nature of mind seems so close to the quick. But the history of science reminds us easily enough that at various stages astronomy, biology, and physics have seemed special targets for concern. There has been continued and explicit discussion of these issues in connection with AI. [Taube, 1961; Weizenbaum, 1976; McCorduck, 1979.] I have not included them in the list of intellectual issues because they do not, in general, seem to affect the course of the science. Where some aspect does seem to do so, as in the issue of helping humans or replacing them, it has been included. However, the broader issue certainly provides a thematic background against which all work goes on in the field, increasing its ambiguity, and the broader issue undoubtedly enters into individual decisions about whether to work in the field and what topics to select.

The second example involves Hubert Dreyfus, who has been a persistent and vocal critic of AI. [Dreyfus, 1972.] He has certainly become an issue for the field; however, this does not necessarily produce an intellectual issue. Dreyfus's central intellectual objection, as I understand him, is that the analysis of the context of human action into discrete elements is doomed to failure. This objection is grounded in phenomenological philosophy. Unfortunately, this appears to be a nonissue as far as AI is concerned. The answers, refutations, and analyses that have been forthcoming to Dreyfus's writings have simply not engaged this issue—which, indeed, would be a novel issue if it were to come to the fore.

The third example involves the imagery controversy, which has been exceedingly lively in cognitive psychology. [Kosslyn, Pinker, Smith, and Shwartz, 1979.] The controversy is over the nature of the representations used by humans in imagining scenes and reasoning about them. There is no doubt about its relevance to AI—the alternatives are a classical dichotomy between propositional (symbolic?) representations and analog ones. Thus, at heart, it is a variant of the issue of analog-versus-digital representation, which has received mention. But for reasons that are quite obscure to me, the imagery issue has received hardly any interest in the AI community, except where that community also participates in cognitive psychology. As things stand at the moment, this would be an issue for cognitive science, but it is not one for AI.

Though enumerating intellectual issues exposes a certain amount of the history of a field, even if only from particular viewpoints, some important parts can be missed. These seem to be endeavors that were noncontroversial or where the controversies were merely of the standard sort—of what progress had been made, what subfields should get resources, and so forth. Thus, work on program synthesis and verification goes unnoticed. Also, the major effort in the 1970s to construct speech-understanding systems is barely noticed. Perhaps this is not a valid point about the basic historical scheme but reflects only the unevenness of my process of generating issues. Certainly, there were issues in speech-recognition research both in the 1960s, when Bell Laboratories decided to abandon speech recognition as an inappropriate task, and in the 1970s, when a substantial effort sponsored by DARPA to construct speech-understanding systems was dominated by AI considerations over speech-science considerations. Perhaps intellectual issues are generated from all scientific efforts in proportion to the number of scientists involved in them (or to their square?); all we need to do is look for them.

## Characteristics of the History

Turning to what is revealed in Table 1, the most striking feature, to me at least, is how many issues there are. Looked at in any fashion—number active at one time (fifteen on average) or total number of issues during AI's quarter-century lifespan (about thirty)—it seems to me like a lot of issues. Unfortunately, similar profiles do not exist for other fields (or I do not know of them). Perhaps the situation in AI is typical, either of all fields at all times or of all fields when they are getting started. In fact, I suspect it is due to the interdisciplinary soup out of which AI emerged. [See my paper "Reflections on the Structure of an Interdiscipline" in this volume.] Many other related fields were being defined during the same post-World-War-II era—cybernetics, operations research, management science, information theory, control theory, pattern recognition, computer science, and general systems theory. Even so, I do not see any easy way of pinning down a correct interpretation of why there are so many issues.

Issues are not independent; they come in clusters, which are coordinated. Researchers tend to fall into two classes, corresponding to one pole or another on all issues in the cluster. Clusters that occur in this history are as follows (where polarities of subissues have been reoriented, if necessary, to make them all line up together, corresponding to the superordinate issue):

- AI versus Cybernetics
  - Symbolic versus continuous systems
  - Problem-solving versus recognition
  - Psychology versus neuroscience
  - Performance versus learning
  - Serial versus parallel
- AI versus Computer Science
  - Symbols versus numbers
  - Heuristics versus algorithms
  - Interpretation versus compilation
  - Replacing versus helping humans
  - Problem-solving versus theorem-proving
- Problem-Solving versus Knowledge Search
  - Heuristics versus epistemology
  - Search versus knowledge
  - Power versus generality
  - Processing versus memory
- Linguistics versus AI and Cognitive Psychology
  - Competence versus performance
  - Syntax versus semantics
- Engineering versus Science
  - Engineering analysis versus simulation
  - Engineering versus science
  - Real versus toy tasks
- Wholes versus Atoms
  - Procedural versus declarative representation
  - Frames versus atoms

A cluster might seem to define a single underlying issue, which can then replace component issues. However, the fact that issues are coordinated does not make them identical. Some scientists can always be found who are aligned in nonstandard patterns. In fact, some of the clusters seem much more consistent than others. Thus, the multiplicity of issues keeps the scientific scene complex, even though, because of clustering, it appears that it should be clear and simple. In fact, many of the groupings are more easily labeled by how they separate fields than by any coherent underlying conceptual issue.

Clustering of issues does seem to be a common occurrence; for instance, a standard advanced text on learning in psychology begins with a list of seven dichotomous issues that characterize learning theories. [Hilgard and Bower, 1948 and 1975, pp. 8–13.] The first three—peripheral versus central, habits versus cognitive structures, and trial-and-error versus insight—form a coordinated cluster that characterizes stimulus/response theories versus cognitive theories (to which could even be added tough-minded versus tender-minded, the contrast William James used to distinguish the two main types of psychologists). One possible source for such coordinated clusters is the attempt to find multiple reasons to distinguish one approach from another. The approach comes first and the issues follow afterward. Then the issues take on an autonomous intellectual life and what starts as rationalization ends up as analysis.

A major role of the issues here seems to be to carve up the total scientific field into disciplines. AI, computer science, logic, cybernetics, pattern recognition, linguistics, and cognitive psychology—all these seem to be discriminated in part by their position on these various issues. The issues, of course, only serve as intermediaries for intellectual positions that derive from many circumstances of history, methodological possibilities, and specific scientific and technical ideas. Still, they seem to summarize a good deal of what keeps the different fields apart, even though the fields have a common scientific domain.

Is the large burst of issues that occurred at the birth of AI just an artifact of my intent to gather issues for AI? If the period just before AI began, say from 1940–1955, were examined carefully, would many more issues be added? The relevant question should probably be taken with respect to some other field as a base. Would a burst like this be found for cybernetics, which started in 1940–1945? My own suspicion is yes, but I have not tried to verify it.

Perhaps then the situation of AI could turn out to be typical. We would find a plethora of issues in any science if we would but look and count; the list from Hilgard and Bower might serve as a positive indicator. However, before rushing to embrace this view, some counterevidence should be examined. An interesting phenomenon in this same postwar period was the emergence of several one-theorem fields. Game theory, information theory, linear programming, and (later) dynamic programming—all had a single strong result around which the field grew.[15] Certainly, each also provided a novel formulation, which amounted to a class of systems to be used to theorize about some field. But initially there was only one striking theorem to justify the entire field. It gave these fields a curious flavor. My personal recollection is that all these fields, while exciting, profound, and (sometimes) controversial, had none of the complexity of issues that we find in Table 1.

## Intellectual Issues and Progress

There is a natural temptation to use the history of intellectual issues to measure progress, once it has been explicitly laid out. It is true that some issues have vanished from the scene, such as symbols versus numbers; that seems, perhaps, like progress. It is also true that other issues seem to recur, such as problem-solving versus recognition; that seems, perhaps, like lack of progress. Neither interpretation is correct, I think. Rather, the progress of science is to be measured by the accumulation of theories, data, and techniques, along with the ability they provide to predict, explain, and control. This story is not to be told in terms of such intellectual issues as populate this paper. It requires attention to the detailed content, assertions, and practice of the science itself. True, at the more aggregate level of the *paradigms* of Kuhn or the *programmes* of Lakatos, whole bodies of theory and data can become irrelevant with a shift in paradigm or programme. But on the scale of the twenty-five years of AI research (1955–1980), the story is one of accumulation and assimilation, not one of shift and abandonment. It is not even one of settling scientific questions for good.

What then is the role of intellectual issues in the progression of science? To echo my earlier disclaimer, I can only conjecture. Intellectual issues seem to me more like generalized motivators. They evoke strong enough passions to provide the springs to action, but they are vague enough so that they do not get in the way of specific work. They can be used to convey a feeling of coherence among investigations in their early stages, before it is known exactly what the investigations will yield.

Evidence for this is that issues do not really go away; they return and return again. Repetition is abundant in Table 1. The model that suggests itself immediately is the spiral—each return constitutes a refined version of the issue. Though the issues are certainly not identical each time, it seems difficult to construe the changes as any sort of progressive refinement; some seem more like wandering (e.g., the serial/parallel issue). A more plausible explanation (to me) is that intellectual issues reflect perennial unanswerable

[15] Another field, general systems theory, also had a single idea around which to build—that there are common laws across all levels of systems from the atomic through cellular through societal through astronomical. But there was no central result available, only the system view, and this field has been markedly less successful than others in its growth and health.

questions about the structure of nature—continuity/discontinuity, stasis/change, essence/accident, autonomy/dependence, and so forth. Whenever in the course of science one of these can be recognized in the ongoing stream of work, an appropriate intellectual issue will be instantiated, to operate as a high-level organizing principle for a while. To be sure, this picture does not capture all that seems to be represented in our population of intellectual issues. But it seems substantially better than viewing science as progressively resolving such issues.

## CONCLUSION

Putting to one side questions about the accuracy of the particular set of issues displayed in Table 1, of what use is a history of a scientific field in terms of intellectual issues? To repeat once more: It cannot substitue for a substantive history in terms of concepts, theories, and data; however, it does seem to capture some of the flavor of the field in an era. It is clearly a component of the paradigm of a field or of research programmes within a field. And, let us confess it, intellectual issues have a certain spiciness about them that makes them fun to talk and write about. Perhaps it is the sense of touching fundamental issues. But perhaps it also echoes Bertrand Russell's famous aphorism that dealing with intellectual issues has all the advantages of theft over honest toil.