

Introduction to the COMTEX Microfiche Edition of the Early MIT Artificial Intelligence Memos

Marvin Minsky

MIT

THESE ARE THE VOYAGES of the MIT Artificial Intelligence Laboratory, and these remarks may help to understand the context of this collection, though in many ways the memoranda speak quite clearly for themselves and my comments are not, in any case, to be regarded as history, for I have written them quite hastily, in much the same spirit of the memos themselves, when it was our strategy in those early days to be unscholarly; we tended to assume, for better or for worse, that everything we did was so likely to be new that there was little need for caution or for reviewing literature or for double-checking anything. As luck would have it, that almost always turned out true.

In any case, my memory is as treacherous as anyone's, and trying to recall those times—even the shapes of spaces where we worked—yields only sharpened images that show the artificial clarity of careful composition; they show “the ways things must have been” as much as “really how they were.” We had so little sense, in those days, of “history in the making” that project after project left no written trace at all. No memos describe the eyeglasses we made with tiny CRTs projecting solid images before the eyes, or the little wireless computer terminal (which was promptly stolen) the complicated automatic “interfaces” we made for the first robotic experiments. When the students made the first “computer games” they were addictive then as now and finally we banned the best one, SPACEWAR, from daytime use with no sense of how historic it was—I think in 1964—to make the first such ban! And no one ever thought to document—or even photograph—our first “Computer Graphic” blazing

galaxies and dazzling worlds of moving shapes, regarded as mere “hacks,” and mainly used to amaze our visitors.

(In fact, because we had so little sense of history, some memos seem to be forever gone; the ones of which I couldn't find a single copy; if anyone out there has one of those, please send a copy we can add to this collection. The other missing memos are because we couldn't get permission to reprint them; in that case send permission, if you can!)

We started the MIT-AI laboratory shortly after John McCarthy and I became Assistant Professors of the faculty of the MIT Mathematics Department. While we had different views about how to understand intelligence, we were both obsessed to find out how a mind could do its “common sense reasoning.” He aimed more toward establishing good logical and mathematical foundations for reasoning, and more toward computer architecture. I pursued questions about heuristic control of problem solving systems, and ideas about how brains might work. I preferred to try to formalize heuristic processes, reasoning by analogy, and theories of pattern recognition, while McCarthy worked to formalize not only reasoning, but also programming semantics and computer methodologies; this led him to some of the first systematic ways to prove theorems about programs.

Ideas like these fused into an atmosphere of mathematical power and heuristic adventure, of unbounded ambition and enthusiasm. The laboratory grew steadily in size, and in ideas from many spheres. It was axiomatic that our work would always have as sound mathematical foundations as possible for its time; this was why we were led to some

of the foundations of the modern theory of computational complexity. But we were equally determined to find how mechanisms could be robust enough also to work on problems that were “ill-defined,” to tolerate diverse varieties of ambiguity

The results speak softly for themselves. I don’t mean to boast, but when computer specialists talk of “artificial intelligence,” they’re rarely conscious of the debts they owe to workers in the AI field for tools and techniques that they use every day in every way. The AI laboratories were streams of basic innovations in “ware” both hard and soft. Few remember how hard it was to use computers in those days; programs prepared by perforating cardboard rectangles were “submitted” as decks of perforated cards to massive scheduling bureaucracies, then passed through hired hands of professional machine operators; it all took hours and even days. In frustration, John McCarthy invented what is still today called “time-sharing,” and soon we had the first installation in which several workers could type programs directly into a computer, at the same time, each getting almost instant answers. (The job of running the machine itself became just one more time-shared “job” and those human operators disappeared.)

To do such things, we had to invent all sorts of “paging” and “protection” schemes that now are standard features of all good computers. Our students developed among the first debugging programs—with cute names like “DDT” that persisted long after that insecticide was banned. Time-sharing’s fast response made it feasible to use the computer itself to help composing text and programs, and this became the normal <modus operandi> of our staff; this reached the public only ten years later, with new names like “word-processing” and “user-friendly.” We convinced a small computer company called Digital Equipment to offer time-sharing on their computers; larger companies took longer and by the time they got the idea, DEC had grown larger than most all of them.

Now some of our students began to graduate from the AILAB, and most went on to research of their own; fifteen of my first twenty students became professors. Some helped build up the research group at Bolt Beranek and Newman, where McCarthy had supervised some of the time-sharing research. Around 1964, McCarthy himself left to start a new AI laboratory at Stanford, which quickly became another world center of AI research; now there were three principal AI laboratories, at MIT, Stanford, and (what is now) Carnegie-Mellon University, where Allen Newell and Herbert Simon pursued AI research they had begun in the middle 1950s. SRI soon became another major AI laboratory and now there are many more, but those first three are still the largest ones.

I had started my own work even earlier, as an undergraduate at Harvard. Fascinated with the ideas of McCulloch, Lettvin, Pitts, and Selfridge, in the great days of early cybernetics, in 1956 I joined Oliver Selfridge’s group at the MIT Lincoln Laboratory, and perhaps derived from him and (in turn) from McCulloch my earliest ideals of how to make

a laboratory. The bold support of our adventure by several others is nowhere fairly documented; Jerry Wiesner and Philip Morse found us our early space and funds, brave mathematicians like Norman Levinson, Witold Hurewicz, Claude Shannon and Norbert Wiener supported our intolerable conceits and, as we grew our junior faculty, heroic chairmen Ted Martin and Peter Elias supported us within MIT’s Mathematics and E.E. departments.

Around 1963 my friend and former teacher of psychology, J.C.R. Licklider, went to Washington D.C. to fund research on computers from the Defense Department’s “Advanced Research Projects Agency.” This led to Project MAC at MIT, which was for us a golden age of financial support; for an entire decade we pursued without compromise the scientific directions we thought were best. Licklider then replaced himself by recent former students of ours, so that our research continued to be sponsored by brilliant scientist-engineers like Larry Roberts and Ivan Sutherland—while playwright-linguist Marvin Denicoff administered the contracts through the Office of Naval Research. At MIT the golden age of Project MAC was masterminded by Robert Fano, a visionary engineer-theorist.

Soon after McCarthy moved to Stanford, Seymour Papert joined the laboratory; mathematician, psychologist and philosopher, fresh from five years work on child development with Piaget in Geneva. We worked so well together that, for a decade, we each could run the laboratory effortlessly, leaving the other to decide what must be done. Co-directing is never having to discuss non-technical matters. (But this was also due to the giant web of tasks assumed by engineer-manager Russell Noftsker, who also could read scientific minds.) We worked especially closely together on developing “intelligent robotics”; our goal was to make machines both to see, and understand that which they saw, enough to make their metal hands do interesting, real jobs. We had to invent many of the first ways to program mechanical hands and electronic eyes; many such details were documented in myriads of “working papers” outside the present series of memos.

Several streams of AILAB research barely appear at all in these memos. We worked for several years on mathematical theories of the then-mysterious “Perceptron learning machine,” but the results were published only in book form. Papert elaborated a great network of ideas about mental development and education, and built a fertile research group around his new computer language LOGO; that work on the theory of education scarcely shows at all in Memos because it had its own series. I simply never got around to document, at all, three years of building a powerful LOGO-based personal computer for real-time animated graphics. And while LOGO research began in the late 1960’s, its ideas entered the mainstream of public education only in the early 1980’s; we had to wait so long for inexpensive personal computers that many of these “new” ideas are older than the children learning them.

As a record of the laboratory’s work, this collection

has other dimensions of incompleteness. The "AI memos" were meant to be informal; little more than half-baked schemes, written down for others (both inside and outside the laboratory) to use, but not deemed finished or significant enough for formal publication. When ideas became "final" documents, e.g., academic Theses, or final Project Reports, they appeared either as MIT "Technical Reports" or as publications in professional journals. Also there were other series of "working papers" that were regarded as entirely internal, for use by colleagues in the laboratory, and not advanced enough to become "AI memos," which often contained program listings, more details of how things actually worked, circuit diagrams, or even charts of where machinery was hidden under section of laboratory floor. But often there were substantial discoveries not documented elsewhere.

But the largest and obscurest dimension of incompleteness was the paradox in which the things most clearly understood are scarcely written down at all; why bother if you can explain it in a few moments? Yet such ideas that people carried only in their heads were often just the most important ones. It was only by the merest chance that I actually wrote the widely influential Memo 306,^{*} because when I first explained the idea everyone said it was obvious. It seemed worth writing down only when I chanced on someone who DIDN'T like it.

In the early 1970's the mood at ARPA changed for a time, perhaps because we hadn't found another brilliant youngster to spend that precious time in Washington. Administering and funding the AILAB became more difficult, and ARPA even tried to influence the content of our proposals. Neither Papert nor I could deal with this but Patrick Winston, one of our most original researchers, could. So we made him Director, and he still is. The memo record shows less frequent Winston papers, once he undertook so much at once, but that written record fails to show the growing influence of his ideas on student, staff, and faculty.

What is "Artificial Intelligence," anyway? It makes no sense to define things not inherently constrained; AI was what one made it be - when none were wise enough to specify the very best thing it could be. It is easier to say what it AI wasn't: it's absolutely not that kind of mish-mush "interdisciplinary" combination that comes and goes in universities. Instead, I see it as a science of its own, growing with increasing coherency, of what processes—it makes no difference that they be embodied in machines or minds—can solve various sorts of problems. For practical purposes we usually tell passers-by this easy definition: "AI concerns performances that a person needs intelligence to do." For instance, when Slagle wrote the SAINT program in 1960, that was "AI," because solving college calculus problems then seemed to need intelligence. However, once Jim Slagle showed us how, such problems somehow no longer seemed to need so much intelligence; in fact it left us wondering why

students take so long to learn to solve those kind of problems.

So, in this sense, the term "intelligence" itself seems only to describe the moving horizon of our growing understanding of how minds might work. Imagine, if you dare, that some super-intelligent extraterrestrial alien were to examine a human brain and understand completely how it works, just as a person understands completely all the gears and ratchets of a wrist-watch. So smart an alien might not consider US intelligent. The idea it might talk to us would seem as strange, to it, as we would think of talking to a clock. I find most people seem un-nerved at the idea of a science aimed at a moving horizon. They say accusingly: "how can you study something that you can't even define." Well, most studies are just so; Biology, the study of "life" is precisely the horizon of our growing understanding of organisms. Once Watson-Crick showed us how gene-strings reproduce, it left us wondering why people took so long to think of such a simple thing—and every scientist knows now that there isn't any real boundary defining "life" except that moving frontier marked by what we understand of physiology.

Indeed, from yet another point of view, I sometimes think of AI as "the current frontier of computer science." (This angers some who call themselves "computer scientists"—but everyone must understand how sensitive must be those colleagues who proclaim their "Science" on their very greeting cards.) Then, in that view, AI is simply finding ways to make computers do the useful things that no one yet knows how to make them do. This lazy comprehensiveness has one annoying side-effect of making AI's cumulative reputation subject to a continual "exponential decay"—wherein each achievement fades away to be credited to some other specialty. Actually, I think this is a great and vitalizing fact; let me explain it by examples.

In AI's early days we were concerned with recognizing patterns of many kinds. Today, "pattern recognition" has become a separate field; it has journals of its own, nor will AI journals accept papers on that subject. Similarly a new field of "symbol manipulation" emerged from AI research efforts like our MACSYMA project, now seen as in the field of "symbolic applied mathematics." Another "spin-off" from AI is the soon-to-be enormous industry of intelligent sensorimotor Robotics. Yet another industry will soon emerge from work in the 1970's on making programs that write programs; this probably will call itself "automatic programming." And fairly soon ideas that have been brewing in AI since the 1960's, on making computers understand significant fragments of natural language, will enter and, I'm sure, soon dominate the main stream of Linguistics. (In the era of these memos, it was the students in AI, almost alone, who carried on the quest for meaningful theories of linguistic processes, when most all other language work was stuck in shallow, syntax-oriented, formalisms.)

In reading through the memos you can almost feel the pace at which "responsiveness" evolved in our computing systems. Only today are such "programming environments" becoming popular "outside," using "new techniques" that stem

*Minsky, M A framework for representing knowledge, MIT AI Memo 306 June 1974

from those earlier, unheralded AI systems that integrated editing, debugging, and compiling within the single systems that evolved across these twenty years. And now, today, much AI work is aimed at representing knowledge in computers; these new techniques are moving into what is called "cognitive psychology," and also into the new industry of building so-called "expert systems."

These memoranda depict the early days in which those youngsters came to MIT from all the world, obsessed and inspired perhaps as much by science fiction fantasy as scientific papers—as much by Asimov, Heinlein and Pohl as by Turing, Shannon, and McCulloch. But the past few years have strangely changed the social world of Artificial Intelligence; now it has hit the proverbial fan. When those first students came to work with us, they had to bring along a special kind of courage and vision, because most authorities did not believe AI was possible at all. Today the AI laboratories suffer raids from greedy industries, and students need a different sort of strength, sometimes before they even graduate, to turn away from salaries quite higher than commanded by department heads. I dread the embrace of that commerce-world, for soon we'll have to watch our students start to hide "proprietary" papers from their friends: we pioneers had but to share, to face each real or imagined enemy. For then it was impossible to "give away a good idea" while now we're listened to, I feel, a lot too much more carefully.

Besides, I find it queer when entering students ask "what attracted you to AI" or "how did you get interested in computers?" To them it seems such things were always there; to us it seems they've barely yet arrived. So now I'd caution students: "are you sure it's good to be so interested in computers? Shouldn't you try to start to work on what will come *after* computers?" Of course I'd just pretend to be surprised when they're surprised, because I haven't yet myself imagined quite what such a thing might be. (Well, nothing like a present-day computer, but probably some sort of active-memory semantic network—and surely made of solid optics or something, because those 2-D "chips" waste too much space and therefore will not last too long.)

So in that one, quite different, sense this set of memos does have a certain island-like complete integrity: it spans the era within which AI attained complete respectability. This doesn't herald any punctuation in research; why, just the problems in clear view could fill another hundred years, with scientific and technological challenges more delicious than ever before. I hope these memos show the wonder and exhilaration of how it was to grow with and in that terribly new, marvelously intricate, and—as it turned out—entirely sound intellectual domain. But I suppose there's really just one way to fully share in such a thing; you have to take some big exhilarating risk and reject almost all old theories, thinking—"how wrong those early AI workers were, how foolish and how careless they must have been to miss some of the best and simplest things"—and then explore some very different path and hope it leads to something very good.

Enter The New ROBOTICS AGE

An age where formerly impossible tasks become everyday reality. The mechanization of many laborious physical and intellectual tasks of industrial, commercial and domestic life is coming about through the technology of intelligent machines. We see innovative applications of microprocessors, simulations, sensor electronics, real time control, effector design, machine design, aerospace robots and autonomous systems.

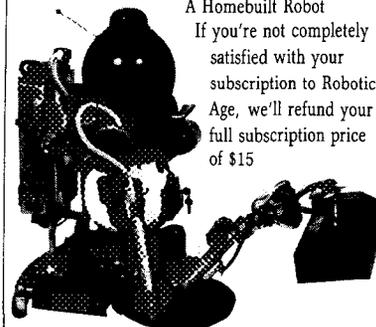
This is the world of Robotics Age, the Journal of Intelligent Machines. Make this world yours. In Robotics Age you'll find topics from the realities of industrial automation to the prospect of personal robotics. Integrating it all is modern software engineering. As illustrations, we show how you can perform practical design explorations with inexpensive personal computers.

Recently published titles include: Avatar, A Homebuilt Robot; The Physics of One-Legged Mobile Robots; Constructing an Intelligent Mobile Platform; An Inexpensive Hand; Natural Language Understanding; Applying Robot Vision To the Real World; Telecommunications Robots. In every issue you'll find features like Patent Probe, New Products, Letters and Design Forums.

Don't miss out. Subscribe today. With your paid subscription accompanied by the coupon below, we'll send you a free reprint of the article Avatar:

A Homebuilt Robot

If you're not completely satisfied with your subscription to Robotics Age, we'll refund your full subscription price of \$15.



meet AVATAR: A Homebuilt Robot

Sign me up for a 6 issue trial subscription
 Send an issue and bill me \$15.00 for 6 in all
 Enclosed is \$15.00 (6 issues plus Avatar reprint)

Name _____

Address _____

Town _____

State _____

Zip _____

Charge My Master Card

Visa

Number _____

Expires _____

Robotics Age Magazine
 P.O. Box 358 - Peterborough NH - 03458
 (603) 924-7136