AI Magazine Volume 5 Number 3 (1984) (© AAAI)

Introduction to the COMTEX Microfiche Edition of Reports on Artificial Intelligence from Carnegie-Mellon University

Allen Newell

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

ORIGINALLY IT WAS *Complex Information Processing.* That was the name Herb Simon and I chose in 1956 to describe the area in which we were working. It didn't take long before it became *Artificial Intelligence [AI].* Coined by John McCarthy, that term has stuck firmly, despite continual grumblings that any other name would be twice as fair (though no grumblings by me; I like the present name). *Complex Information Processing* lives on now only in the title of the *CIP Working Papers*, a series started by Herb Simon in 1956 and still accumulating entries (to 447). However, from about 1965 much of the work on artificial intelligence that was not related to psychology began to appear in technical reports of the Computer Science Department. These reports, never part of a coherent numbered series until 1978, proliferated in all directions. Starting in the early 1970s (no one can recall exactly when), they did become the subject of a general mailing and thus began to form what everyone thinks of as the CMU Computer Science Technical Reports.

A famous aphorism of Lord Kelvin has it that, if you can't measure something, you really don't know much about it. I don't know how appropriate the aphorism is for science, its presumed target. In fact, much of computer science is devoted to knowing things that are without numerical measure. But the aphorism does seem to apply to report series. Without a well-numbered series, with clerks slaved to mark off each publication, one by one, through the ages, it is only a question of time before no one knows the assemblage or its boundaries. Marvin Minsky, in his corresponding introduction to the MIT Artificial Intelligence Memos (Minsky, 1983), observes how imperfect its coverage was and how much was missing. But at least the integers assure us that it is incomplete. With the CMU collection there is only the prototypic concept of a series and a few archetypical papers that are surely part of it. (Where the boundaries lie, what is to be included in, and what out is a task for present scholars). It is no doubt a worthy task, and several people at CMU have entered into it with the usual spirit and perseverance. A suitable collection has been assembled which captures much of the research done in artificial intelligence over almost thirty years. It is available as a permanent and useful archive within the Comtex series.

Of course, there are *reasons* why there is no numbered series of CMU artificial intelligence technical reports. These reasons tell something of the story of how artificial intelligence research developed at CMU and provide another small perspective on the early history of artificial intelligence.

CMU has never had an Artificial Intelligence Laboratory. The research on artificial intelligence has always been embedded in other institutional frames at CMU. In the earliest days (the second half of the 1950s), we were located in the Graduate School of Industrial Administration [GSIA]. Herb Simon was an Associate Dean there, one of the small band who had started GSIA just a few years earlier (1949). GSIA was by then already well launched on its (successful) attempt to revolutionize business schools, moving them from the then-dominant case-study method of the Harvard Business School to the now-dominant blend of social science, economics and operations research (from practice-oriented to applied-science oriented.) There was plenty of revolution

I thank John McDermott, Alan Perlis, Raj Reddy, and Herbert Simon for comments on an earlier draft. I would also like to express my appreciation to Sylvia Hoy for her efforts in getting the collection together

to go around, and rather than becoming something separate, work on complex information processing simply became one more aspect of the new look in the science of decision making. All the early work in AI occurred as PhD's in Industrial Administration—Ed Feigenbaum's (1959), Julian Feldman's (1959), Fred Tonge's (1960), Geoff Clarkson's (1961), Bob Lindsay's (1961), and indeed my own (1957). The CIP Working Papers record most of these theses, but not many working papers that did not eventually see the light of published day.

Actually, the CIP Working Papers do not contain all the early work on artificial intelligence at Carnegie. The research was originally conducted by what became known as the Rand-Carnegie group. J. C. (Cliff) Shaw and I were both at the Rand Corporation (although I was physically located at Carnegie). Thus, in the very first years before students in GSIA became much involved, much of the research appeared as Rand Corporation's *Research Memoranda* and *Papers*. As a proper research organization, of course, they had a proper technical report series. For instance, the first report on the Logic Theory Machine was published as Rand Paper P-850 in May 1956 (Newell & Simon, 1956). Gradually, after my own organizational affiliation shifted to Carnegie in 1961 and Cliff's work focused on timesharing (the **Joss** system, Shaw, 1964), use of the Rand publication series trailed off, although it did not cease entirely until 1963 (Newell, 1963).

It is interesting to reflect how strongly intellectual fields are influenced by the institutional frameworks in which they are placed. Because the early efforts took place in GSIA, the early directions of artificial intelligence at Carnegie were in management science. Assembly line balancing (Tonge, 1960), portfolio selection (Clarkson, 1961), job shop scheduling (Gere, 1962), and warehouse location (Kuehn & Hamburger, 1963) were all subject to study. The term *heuristic programming* came into use, in homology to *linear programming* and *dynamic programming*, to describe AI programs in management science. These two earlier uses of the term *programming* were newly coined in the post-World-War-II era, and they owed nothing to the programming of computers. They traced their roots to George Dantzig's initial efforts in the Air Force Office of Strategic Programs in the late 1940s to develop schedules (*i.e.,* programs) of expenditure and acquisition (Dantzig, 1951). This developed into linear programming. When Richard Bellman came along in the mid-1950s he called his scheme *dynamic programming* (Bellman, 1957). With heuristic programming, of course, the double meaning of both a type of management science programming and a type of computer programming was manifest.

When, shortly thereafter (to be recounted below), the work in AI shifted out of GSIA to other organizational bases in computer science and psychology, the growth in heuristic programming for management science slowed. It did not stop, of course, but continued at many institutions as individual efforts, as a part of management science and operations research. But it moved out of GSIA before other management-science institutional bases had grown strong enough to take intellectual leadership. Thus, it became only a minor theme within management science, although that may now be changing with the spread of interest in expert systems. Interestingly, in the light of this current wave of interest in AI applications, the AI research in management science always had an orientation to real tasks and to applications. However, the emphasis was not on being knowledge intensive *per se*, but on the complexity of the problem formulation and on being heuristic; *i.e.,* finding good, not optimal, solutions; just those aspects that tended to distinguish heuristic programming from mathematical programming. It should be noted that one corner of the work in GSIA had a strong psychological flavor, stemming from the general interest in management science in decision making in the individual (Feigenbaum, 1959; Feldman, 1959; Clarkson, 1961).

In those days, of course, there were no Computer Science Departments. They did not emerge anywhere until 1964 and not at Carnegie until 1965. There was, however, an interim organization that weaned AI research from GSIA. Starting in 1961, Carnegie had a *Systems and Communication Sciences [S&CS] Interdisciplinary Graduate Program.* In those days, almost everyone believed that *communication sciences* was the magic term to cover the burgeoning flux of intellectual activity around control theory, communication theory, linguistics, computers, cybernetics and systems theory. The University of Michigan had led the way, if I recall, with a Communication Sciences interdisciplinary program and MIT had followed suit with a Center for Communication Sciences. Even then, it is interesting to note, institutions in the United States could not bring themselves to call the field *cybernetics*, although we had many discussions at Carnegie about adopting that name. The personal stamp of Norbert Weiner seemed too much upon it. Ultimately, of course, the computer had its own way and carved out a new disciplinary niche. This rent a large hole in the middle of the communication sciences umbrella, and sealed its fate to providing only momentary intellectual cover.

The S&CS program at Carnegie was built around GSIA and the departments of Mathematics, Psychology and Electrical Engineering.[1] Its official leader was Alan Perlis and its physical locus was the Computation Center (of which Alan was Director). The other principals, in addition to Herb Simon, were Peter Andrews (Mathematics, in logic and theorem proving), Bert Green (then head of the Psychology Department, in cognitive simulation and psychometrics), Lee Gregg (Psychology, in learning), Abe Lavi (Electrical Engineering, in systems) and Bill Pierce (Electrical Engineering, in information theory). This interdisciplinary matrix became the home of research in AI and much else besides.

The S&CS program never initiated a report series. I don't think it ever occurred to anyone to do it. Perhaps there was simply too much eclecticism; it was not imaginable

[1] Carnegie has never had a Linguistics Department, so these four departments covered all that fit within the notion of communications and systems

that the varied work of all the participants was relevant to a single audience. Perhaps, for everyone involved, the S&CS effort was not so much a new beginning in our research lives as simply an organizational frame for what we were already doing. Its main effect was perceived, if my recollection is correct, to provide a more convenient framework for graduate education.

If that was our perception, it turned out to be only a partial view. In 1962, J. C. R. Licklider took leave from MIT to set up the Information Processing Techniques Office [IPTO] of the Advanced Research Projects Agency of the Department of Defense [DARPA]. The history of that effort needs to be set out in some appropriate place, for as both Marvin Minsky and Bruce Buchanan emphasize in their respective introductions to the MIT and Stanford report series (Minsky, 1983; Buchanan, 1983), the DARPA support of AI and computer science is a remarkable story of the nurturing of a new scientific field. Not only with MIT, Stanford and CMU, which are now seen as the main DARPA-supported university computer-science research environments, but with other universities as well (UCLA, Illinois and Utah come immediately to mind), DARPA began to build a remarkable nationwide research community. We received our own first research contract in 1962, with an open invitation to build excellence in information processing in whatever fashion we thought best. This broad contract, given to Alan and myself as principal investigators, found an obvious and immediate institutional home in the budding S&CS program. (The DARPA effort, or anything similar, had not been in our wildest imaginings during the formation of the S&CS program.) No further institutional arrangements were needed and none were forthcoming. Certainly, no new impetus was provided for initiating a report series. We simply continued what we were doing, only now with adequate support and opportunities for growth.

As noted, in 1965 this institutional framework transmuted into a Computer Science Department. It would be satisfying to see this as a major institutional event, arising out of strong forces and needs. Ideologically, perhaps, a case can be made for this. I remember it mostly as Alan Perlis's doing; he remembers it mostly as the students' doing. Early on, Alan became convinced that the computer was destined to bring forth a scientific field. The students couldn't understand why, if the computer was the center of scientific and intellectual concern, they had to decide whether to be a mathematician or an electrical engineer or whatever. Alan finally decided that it was appropriate to shift our identity and it happened. It was an event of the mind, to accommodate what was happening in the world of ideas. Institutionally, on the local scene, there was no need for it and indeed institutionally nothing changed. The thirty-odd S&CS graduate students that had clustered around the Computation Center and were supported by DARPA research funds simply became computer science students and continued to increase. We all continued with our research, on AI and in other areas. A computer science graduate curriculum emerged, but with a strong systems flavor; for example, operations research was included in the qualifying examinations. Importantly, the department was purely a graduate one. This also was in part ideological. We believed that computer science was still too young to justify undergraduate degrees; but in part, it was just a continuation of the fact that S&CS was a graduate program. Its effect was to keep education and research completely identified, so that no need was ever felt to create separate institutional structures for research.

AI research continued wherever it had been going on, and organizational labels just fell as they happened. Mostly, research occurred in the Computer Science Department and the Psychology Department, with a little in Mathematics (Peter Andrew's work on theorem proving) and a little in GSIA, continuing from the initial concentration there. Herb moved from GSIA to the Psychology Department in 1968. Broad program support was the order of the day in the 1960s. Centered in Computer Science, of course, was the DARPA support. But also, centered in Psychology, was a program grant from the National Institute of Mental Health [NIMH](which also began in 1962, with Bert Green and myself as principal investigators, although Herb soon took it over). The National Institutes of Health (of which NIMH was then a part) was deeply committed to exploiting the computer and all its implications for the medical sciences. Thus, support flowed from several sources, to wherever it was appropriate, and research simply continued in its individual ways.

The eclecticism of the times and of the environment at Carnegie, and the lack of strong research organizations and projects, seems a little startling in retrospect. The distinction between engineering-oriented AI and psychologically-oriented AI is well established in the current mind. It was already alive and well in the 1950s (*e.g.,* see the introduction to Samuel, 1959).[2] But no such distinction played a significant role at Carnegie, except when we talked to the outside world. To be sure, the more engineering-oriented AI research tended to be done in Computer Science and the more psychologically-oriented tended to be done in Psychology. But that was just a diffusion effect. The heritage of AI research in management science at Carnegie, which could be pigeon-holed neither as engineering nor as psychology, assured that no dichotomous description was apt.

The distinction between computer science and artificial intelligence also played only a negligible role at Carnegie. Internally, the Computer Science Department was organizationally seamless. There were no academic divisions and no projects. The DARPA funding was used for the entire spectrum of research—time-sharing, programming languages, networking, multi-processors, AI, theoretical computer science, computational linguistics. This was in marked contrast with the situation at our sister institutions, Stan-

[2]The term *Cognitive Science* as the umbrella term to cover AI research with a human orientation, either psychological, linguistic or anthropological, doesn't emerge until about 1977, with the journal, *Cognitive Science*

ford and MIT. At Stanford, DARPA funding supported the AI Laboratory exclusively, and the Laboratory was even physically distant from the rest of Computer Science. At MIT, DARPA funding supported both the MULTICS project (now transformed into the Laboratory for Computer Science) and the Artificial Intelligence Laboratory, cheek by jowl but separate. Both the Stanford and MIT AI Labs had single strong leaders and, protected by organizational boundaries, each grew tight research cultures. At Carnegie, the leadership was always shared, originally between Alan Perlis, Herb Simon and myself, and it always spread across organizational boundaries. Thus the research culture was more diffuse.

Even with the formation of the Computer Science Department, no well- organized technical report series emerged, although Computer Science did finally begin to put covers on reports and distribute them. In fact, I recall strong feelings against being counted, and thus boxed in. We saw no reason why anyone should be able to encompass in one view all the research our environment we were engaged in. Why should we be forced by the existence of obligatory lists to say all that we were doing? An echo of that resolve still exists, the *CMU Computer Science Research Review.* We started it in 1966, determined neither to list nor totally order the environment. The opening introduction of the first issue preserves the flavor:

> This document tells something about research in information processing at Carnegie Institute of Technology in 1966. It tries to say it mainly by a series of essays, written by some of us in the environment, that reveal aspects of computer science of concern here. Although we have included a certain amount of descriptive material, *i.e.*, listings of people, reports, and so on—we have avoided the long compilation of small paragraphs of progress, common to most progress reports. Such compilations have their uses, but mostly they just create a fiction They present the appearance of a neat organization of research—here is what is going on at Carnegie (or wherever). But research cannot be so packaged, and the picture of a social process under control is largely spurious. In the main this is because research efforts are related to their ultimate goals by a strong bond of hope, as well as by a weak chain of rationality. In the idiom of problem solving programs, one has at best tests to avoid foolishness No reliable differences can be had between a current state of knowledge and a desired one. To be sure, one must move forward and explore. So one picks a goal; one decides to build a new programming language, or to prove that a program does what it claims. But the goal itself is only a surrogate, only a means to an end. There will be no difficulty recognizing the end: the new technique; the new insight into the nature of information processing; the new whatever; each will be clear enough when it occurs (at least to a small subset of the field). But these final results often bear only a tangential relation to the initial surrogate goals.
>
> What, then, can be said about progress? Certainly the scientific papers that have been produced should be put forward. They represent science in units that seem appropriate to the scientists. The public and social character of science says that each piece of work shall be communicated to the field in a packaging of the scientist's own choosing. But beyond that, perhaps, a communication whose degree of precision matches the reality is most appropriate. That is what this report attempts to be (Introduction, 1966)

Although this publication served as a general communication to all who were interested in computer science at Carnegie, it also served for several years as our official progress report to our sponsors at DARPA. It was our alternative to the famous MIT Research Laboratory for Electronics *Quarterly Research Progress Reports*, which included summary statements of each area of work. Useful though these RLE reports have been to the rest of the world, they always seemed to us a real burden to the researchers. No report series for us!

The flux of change ever washes over the affairs of men. Carnegie is no exception. In 1967, Carnegie Institute of Technology became Carnegie-Mellon University. In 1970 the integration of the Computer Science Department and the University Computation Center (with Alan Perlis as the head of both) came to an end. We split up and went our separate ways, with the Computer Center continuing with the IBM 360/67 and Computer Science building its own PDP10 facilities, thus joining the MIT and Stanford AI Laboratories in common equipment, just as the ARPANET came into being. Perhaps most important of all, Raj Reddy joined the Computer Science Department in 1969, coming from Stanford, where he was a member of the AI Lab there. With Raj came a major change in the AI research done at Carnegie, moving us into machine perception, first speech, then vision, and ultimately into robotics. Up to this time we had focused entirely at the symbolic level, following out a research destiny that, although broad, was a projection of the fields of management science, cognitive psychology and programming. We had ignored the forays of all the other major AI research centers—MIT, Stanford and SRI—into robotics in the late 1960s. Raj was the first scientist trained in another AI Lab to come to CMU, and he brought with him an intense concern with perception, manipulation and motion. Our AI research was transformed, and Carnegie became a major participant in the DARPA Speech Understanding Systems Research program in the mid-1970s (Newell, Barnett, Forgie, Green, Klatt, Licklider, Munson, Reddy & Woods, 1973). In 1980, again due in large part to Raj, a new research organization, the CMU Robotics Institute, was formed, whose research program spread out to a goodly fraction of all the disciplines in the university, including engineering, statistics, mathematics, management, English and social science. As for the present, we are awash, as are all centers of AI research, with the current enthusiasms for applying AI and we have a substantial effort in expert systems.

Through all this flux, the organization of AI research has remained essentially constant. Its primary residence has remained embedded within the rest of computer science, simply as one aspect of the departmental structure. Most of the psychologically oriented aspects have remained tied into the Psychology Department, providing a continuous bridge to mainstream cognitive psychology. And, although the environment has now grown several-fold over that of the 1960s and early 1970s, the AI research remains more like an amorphous mass than a well-ordered organization. The Speech Understanding System Research effort, along with the construction of the C.mmp multiprocessor (Wulf, Levin & Harbison, 1980), did indeed initiate large research projects to the Computer Science Department in the early 1970s. But our projects have generally proved to be without organizational walls. Similarly, the Robotics Institute has certainly provided a more structured world for its research, along classical research-management lines. At least it is so to surface view. But below the surface, it retains many aspects of the old amorphous organization and especially so in its relationships with Computer Science, with whom it shares research facilities, offices, common researchers, and a common view of the world. The organizational boundaries are not without their effect, of course, but maybe the transboundary character of the research environment since its beginning accounts for how small these effects really are so far.

However, some things do at last change. Numbered technical reports have finally arrived in a true report series. The Computer Science Department started numbering in 1978 and the Robotics Institute had a series from the beginning. But of course, the Computer Science series is not for AI research alone, but for all of computer science. Hence, only a modest fraction of its entries represent AI research. Conversely, the AI research in Psychology is not covered in the Computer Science series, though it continues to remain as deeply involved as ever through the work of John Anderson, Pat Carpenter, the late Bill Chase, Dick Hayes, Marcel Just, David Klahr, and Jill Larkin, in addition to Herb. Thus, for another while at least, we avoid being completely charted.

Which brings us back to the present Comtex series. What should go into it? In the spirit of the series, it has seemed to us it should contain most of the material that has had a strong bearing on AI research from whatever source it came, starting with the Rand and the early CIP Working Paper series, and moving though the unnumbered years. We have included the AI-oriented pieces of work in cognitive psychology from the CIP series, right up to the present, but have dropped out work that is more purely psychological, especially when it is readily available in the archival literature. This captures some of the work of the psychologists above, though not all (for example, not Anderson's work on cognitive skill or architectures for cognition, Anderson, 1983). Finally, we have included the work from the Robotics Institute, which is, of course, quite recent.

It seems to me that this archive, with all its messiness and willful imperfections, should be dedicated to Carnegie, as an institutional environment that for almost thirty years has put substance ahead of form.

References

Anderson, J R (1983) *The Architecture of Cognition* Cambridge, MA: Harvard University Press.

Bellman, R. (1957) *Dynamic Programming.* Princeton, NJ: Princeton University Press.

Buchanan, B. G. (1983) Introduction to the Comtex microfiche edition of Memos from the Stanford Artificial Intelligence Laboratory. *AI Magazine*, Vol 4, No.4, 37-42.

Clarkson, G P. E (1961) Portfolio Selection: A simulation of trust investment. Doctoral dissertation, Carnegie Institute of Technology.

Dantzig, G. B. (1951) The programming of interdependent activities: Mathematical models In Koopmans, T C. (Ed.), *Activity Analysis in Production and Allocation: Proceedings of a conference.* New York: John Wiley and Sons.

Feigenbaum, E. A. (1959) An Information Processing Theory of Verbal Learning. Doctoral dissertation, Carnegie Institute of Technology

Feldman, Julian (1959) An Analysis of Predictive Behavior in a Two-choice Situation. Doctoral dissertation, Carnegie Institute of Technology.

Gere, W. S., Jr. (1962) Heuristics in Job Shop Scheduling Doctoral dissertation, Carnegie Institute of Technology.

Introduction (1966) Computer Science Research Review. Computer Science Department, Carnegie Institute of Technology.

Kuehn, A. A. & Hamburger, M. J. (1963) A heuristic program for locating warehouses. *Management Science* 9, 643-666

Lindsay, R. K. (1961) The Reading Machine Problem Doctoral dissertation, Carnegie Institute of Technology.

Minsky, M. (1983) Introduction to the Comtex microfiche edition of the early MIT Artificial Intelligence Memos. *AI Magazine*, Vol. 4, No 1: 19-22.

Newell, A. (1957) Information Processing: A new technique for the behavioral sciences Doctoral dissertation, Carnegie Institute of Technology.

Newell, A. (1963) A Guide to the General Problem Solvers Program-2-2 (Technical Report 3337). The Rand Corporation February.

Newell, A & Simon, H. A (1956) Current Developments in Complex Information Processing (Technical Report 850). *RAND,* May

Newell, A., Barnett, J., Forgie, J W., Green, C., Klatt, D., Licklider, J. C. R., Munson, J , Reddy, D R & Woods, W A. (1973). Speech Understanding Systems: Final Report of a Study Group. North-Holland/American Elsevier (Reprint of original 1971 study).

Samuel, A L. (1959) Some studies in machine learning using the game of checkers *Journal of Research and Development.* 3, 210-229.

Shaw, J. C. (1964) JOSS: A Designer's View of an Experimental On-Line Computing System (Technical Report 2922) The Rand Corporation, August.

Tonge, F M (1960) A Heuristic Program for Assembly Line Balancing. Doctoral dissertation, Carnegie Institute of Technology.

Wulf, W. A., Levin, R. & Harbison, S (1980) *Hydra/C.mmp: An experimental computer system.* New York: McGraw-Hill