

AI and HCI: Two Fields Divided by a Common Focus

Jonathan Grudin

■ Although AI and HCI explore computing and intelligent behavior and the fields have seen some crossover, until recently there was not very much. This article outlines a history of the fields that identifies some of the forces that kept the fields at arm's length. AI was generally marked by a very ambitious, long-term vision requiring expensive systems, although the term was rarely envisioned as being as long as it proved to be, whereas HCI focused more on innovation and improvement of widely used hardware within a short time scale. These differences led to different priorities, methods, and assessment approaches. A consequence was competition for resources, with HCI flourishing in AI winters and moving more slowly when AI was in favor. The situation today is much more promising, in part because of platform convergence: AI can be exploited on widely used systems.

AI and human-computer interaction (HCI) are converging. “Usable AI” conference events in 2008 and 2009 preceded this special issue, and ACM will launch a widely-supported *Transactions on Interactive Intelligent Systems*. AI techniques are in the toolset of more and more HCI researchers, and applications of machine learning are increasingly visible in the HCI literature. Other maturing AI technologies seek input from the HCI community.

The two fields have met under shared tents for some time, notably within *International Journal of Man-Machine Studies* (subsequently *International Journal of Human-Computer Studies*) and at the Intelligent User Interface conferences cosponsored by ACM's Special Interest Groups on Computer-Human Interaction (SIGCHI) and Artificial Intelligence (SIGART). But little of this research has flowed back to the major AI and HCI conferences and journals. In this article, I describe some research that has bridged the fields, but contact has been sporadic.

Logically, they could have been closer. Both explore the nexus of computing and intelligent behavior. Both claim Allen Newell and Herb Simon as founding figures. Working over the years as an HCI person in AI groups at Wang Laboratories, MIT, MCC, and Microsoft, and alongside AI faculty at Aarhus University and the University of California, Irvine, I was puzzled by the separation.

The introduction to this special issue notes the different “monocular views” of interaction with intelligent systems. AI focused on devising better algorithms, HCI on how to improve the use of existing algorithms. AI originated in mathematics and engineering, HCI in psychology. But half a century is enough time to spawn a hybrid or synthesis had forces not pushed the two fields apart.

This article focuses on HCI as represented by SIGCHI, its CHI conference, and the journals *Human Computer Interaction*, *ACM Transactions on Computer-Human Interaction*, and *Interacting with Computers*. Other threads of HCI research appear in human factors and ergonomics, management information systems, and information science. They are part of the story, but CHI research often coexists with AI in computer science departments and curricula, so the paucity of interdisciplinarity there is most striking.

Unlike science, history is more about adopting a perspective—deciding what to include and what to omit—than about being right or wrong. This article is not a conceptual history that tracks the evolution of usable AI. It is a social history, focusing on the forces affecting interdisciplinary work spanning the two fields. It is organized by the eras of AI research, periods where it flourished and periods of wintry neglect.

Relevant publications, including historical sources mentioned and quoted below, are listed in the “Further Reading” section to this article, which also provides examples of papers spanning AI and HCI from the 1970s through 1990s.

Overview of Forces That Kept AI and HCI Apart

AI set out to devise an artificial rival to human intelligence. HCI (as embodied by CHI) has focused on improving applications as they approach widespread use. These markedly different goals led to distinct priorities, methods, and assessment criteria.

CHI has drawn participation and inspiration from companies that develop mass market products, looking for short-term payoffs. AI, strongly focused on future possibilities, tolerated very slow progress. Its principal sponsors and customers, governments, could take a long view, though winters set in when patience ran out.

CHI has focused on technologies with a price point that promises widespread availability, whereas past AI research required expensive mainframe and workstation platforms.

These differences led to a direct tension: AI and HCI competed for intellectual and economic resources. When AI was ascendant, HCI languished; during AI “winters,” HCI thrived.

Rapidly declining costs for storage, processing, and network access have changed this picture dramatically. Today, AI researchers with somewhat scaled-down ambitions work on widely available systems and appreciate the need for greater usability. HCI researchers adopt AI techniques and concerns as part of their everyday work. Greater interaction across the two fields is all but inevitable.

Now let’s see how it unfolded.

The 1950s: AI Origins and the Beginning of a Rollercoaster Ride

The potential of computation was established by its role in code breaking during World War II. After the war, governments funded the building of expensive computers at a few universities. How would they be used? In 1949, a leading code breaker, the British mathematician and logician Alan Turing, created a sensation by writing in the *London Times*, “I do not see why [the computer] should not enter any one of the fields normally covered by the human intellect, and eventually compete on equal terms. I do not think you can even draw the line about sonnets, though the comparison is perhaps a little bit unfair because a sonnet written by a machine will be better appreciated by another machine.” Published a year later were Turing’s “Computing Machinery and Intelligence,” Claude Shannon’s “Programming a Computer for Playing Chess,” and Isaac Asimov’s *I, Robot*, exploring his three laws of robotics. (See the Further Reading section for specific references to works mentioned and quoted in this article.)

The Macy Foundation sponsored a series of conferences in New York in the late 1940s and early 1950s that brought together leading mathematicians, psychologists, and social scientists. Among the topics discussed was cybernetics, defined by one participant, Norbert Wiener, as the study of control and communication in the animal and the machine. Neural network models (initially called “nervous nets”) were also discussed.

The term *artificial intelligence* first appeared in the call for participation in a 1956 workshop written by American mathematician and logician John McCarthy. The workshop participants’ optimistic forecasts attracted considerable attention. When they collided with reality, a pattern was established that was to play out repeatedly. In the premier issue of *Journal of Evolution and Technology* Hans Moravec wrote:

“In the 1950s, the pioneers of AI viewed computers as locomotives of thought, which might outperform humans in higher mental work as prodigious as they outperformed them in arithmetic, if they were harnessed to the right programs. . . . By 1960 the unspectacular performance of the first reasoning and translation programs had taken the bloom off the rose.”

The Early 1960s: A Lull in AI Coincides with an HCI Breakthrough

Part of the pattern is that HCI thrived on resources that were freed when interest in AI declined. In 1960, with the bloom off the AI rose, the managers of MIT’s Lincoln Labs looked for new uses for their

massive (for the time) government-funded TX-0 and TX-2 computers. Granted access to the machines, a newly arrived graduate student, Ivan Sutherland, constructed Sketchpad. Sketchpad introduced many interface concepts, including the first graphical user interface. Although Grace Hopper, Brian Shackel, and others were engaged in HCI-related work by then, Sutherland's accomplishments had the most profound impact on the research thread that shaped HCI as it is today.

The Mid-1960s to the Mid-1970s: AI Takes Shape as a Major Research Field

This first AI downturn was so short-lived that it was more of a cold front passing through than a full-fledged winter. The Soviet launch of the Sputnik satellite in 1958 refocused attention on scientific research. In 1960, J. C. R. Licklider published an influential essay titled "Man-Computer Symbiosis," which defined a major role for artificial intelligence in exploiting computers. Licklider became director of ARPA's Information Processing Techniques Office (IPTO) in 1962. AI subsequently received extraordinary support. For example, MIT's Project Mac funding began at \$14 million (in today's U.S. dollars) for 1963 and rose to over \$20 million annually in the late 1960s. AI research also thrived abroad, especially in the UK. IPTO steadily increased the number of AI laboratories that it funded. This gave AI researchers financial independence in their departments and established AI as a field.

HCI slowly took shape in the shadow of AI visions that precluded joint efforts.

Ambitious Visions

Echoing Alan Turing's remarks, Nobel laureate and AI pioneer Herb Simon wrote in 1960, "Machines will be capable, within twenty years, of doing any work that a man can do." Five years later, I. J. Good, an Oxford mathematician and former World War II code breaker alongside Turing, wrote in a prominent periodical, "the survival of man depends on the early construction of an ultraintelligent machine" that "could design even better machines; there would then unquestionably be an 'intelligence explosion,' and the intelligence of man would be left far behind." Good sketched an architecture for such a machine and concluded that it was "more probable than not that, within the twentieth century, an ultraintelligent machine will be built and that it will be the last invention that man need make."

In 1970, Nicholas Negroponte of MIT argued compellingly that for machines to understand the context in which they operate, they must be able to understand speech. He concluded that comput-

ers would be dangerous and unreliable without speech understanding. Funding flowed copiously to language-processing research.

Forty years later, machines can't do what we do and speech technology is far short of comprehending context at human skill levels. Fortunately, there was a hole in Negroponte's logic. Computers are not dangerous as long as they remain tools controlled by people who understand the context. At the time, though, AI researchers believed that machines would soon be autonomous and intelligent, competing with people on equal terms. Negroponte's colleague Marvin Minsky, working on Project Mac, was quoted in *Life* magazine in 1970:

"In from three to eight years we will have a machine with the general intelligence of an average human being. I mean a machine that will be able to read Shakespeare, grease a car, play office politics, tell a joke, have a fight. At that point the machine will begin to educate itself with fantastic speed. In a few months it will be at genius level and a few months after that its powers will be incalculable."

Other leading AI researchers agreed that a "Renaissance machine" would arrive by 1985. The article ended by quoting Ross Quillian:

"I hope that man and these ultimate machines will be able to collaborate without conflict. But if they can't, we may be forced to choose sides. And if it comes to choice, I know what mine will be. My loyalties go to intelligent life, no matter in what medium it may arise."

The point is not to embarrass people for things said long ago; it is to understand the context of the time and the consequences of such statements. An anxious public was well aware that the world had barely avoided a devastating thermonuclear war during the Cuban missile crisis of 1962. Leaders seemed powerless to defuse the Cold War. AI promised a way out. I was a student in 1970. I read the *Life* magazine article and found the vision reassuring, convinced (perhaps by Asimov, whom I had also read) that a chilling showdown would not materialize.

AI researchers inspired by this vision had little use for the era's HCI focus: improving screen layouts, command names, text editors, and the primitive graphical user interfaces found in a few laboratories. If they were correct, intelligent systems would soon solve the world's problems along with any remaining interaction design problems—assuming that interacting with humans was a priority for intelligent machines.

Two Conceptions of Human Intelligence

The AI founders believed that artificial intelligence was on the verge of rivaling human intelligence, but how did they define human intelligence? Did they look to psychology for help? Not according to McCarthy, who wrote:

“As suggested by the term ‘artificial intelligence’ we weren’t considering human behavior except as a clue to possible effective ways of doing tasks. The only participants who studied human behavior were Newell and Simon. [The goal] was to get away from studying human behavior and consider the computer as a tool for solving certain classes of problems. Thus AI was created as a branch of computer science and not as a branch of psychology.”

What is intelligence to a mathematician or logician? Turing, McCarthy, Minsky, Good, and other AI researchers began as mathematicians. Many were logicians. They worked in mathematics departments prior to the emergence of computer science departments in the mid-1960s. Mathematics is a powerful system, much of it constructible from a small set of axioms by repeated application of a small number of rules. Not surprisingly, theorem proving was a favored topic in early AI research. Chess and checkers also involve the repeated application of a small set of rules to a small set of objects, and board games were another common AI research topic. I. J. Good argued that a first step toward an ultraintelligent machine would be one that could play a version of chess in which the first row pieces are randomly placed (though he insisted that each side’s bishops be constrained to different colors).

If intelligence is the accurate application of well-defined rules to a small set of symbolic objects, then a computer is perfectly poised for intelligence. Tireless symbolic processing? No problem!

The recurring conviction that computers will take over their own education and quickly surpass human intelligence reveals a particular view of what intelligence is. Mathematics can be complex, but ambiguity and imprecise definition have no place (*pace* Gödel and Heisenberg). To a logician, ambiguity suggests error, not intelligence. McCarthy’s goal of setting out to “solve certain classes of problems” was fine. But when skill at solving those specific problems was equated with intelligence and seen as a small step from human intelligence, wild predictions resulted.

Human intelligence, and HCI, must address ambiguity and imprecise formulations. We spend our lives overcoming our nervous system’s skill at abstraction and consistency: learning when to avoid being logical while still being right is a complicated process. A central tenet of HCI is that we cannot design interaction of any complexity by the application of formal rules. It is always necessary to try interfaces out with people, preferably early and often.

McCarthy identified Newell and Simon as the only early AI researchers focused on human psychology. However, even Newell and Simon wrote in 1957 “that within ten years, most theories in psychology will take the form of computer pro-

grams, or of qualitative statements about the characteristics of computer programs.” This profound error shows that at the time, equating human intelligence to logical processes was pervasive across AI research.

Overconfidence in the power of rationality was widespread in the 1960s and early 1970s. It dominated approaches to decision making in business and in the government that funded AI efforts. Robert McNamara and his team of “whiz kids” arrived at the Pentagon in 1961 with computers, statistics, and unbounded faith in rationality. Some of their predictions fared no better than those of their AI contemporaries.

Ironically, central to establishing and funding this AI research was a psychologist who was not convinced by the vision. Citing a 1960 Air Force study that predicted that intelligent machines might take 20 years to arrive, J. C. R. Licklider astutely noted that in this interval HCI would be useful: “That would leave, say, five years to develop man-computer symbiosis and 15 years to use it. The 15 may be 10 or 500, but those years should be intellectually the most creative and exciting in the history of mankind.” Ten to five hundred years represents breathtaking uncertainty; suffice it to say that recipients of Licklider’s funding were on the optimistic end of this spectrum.

The Focus of HCI during an AI Decade

HCI had difficulty competing for funds and students against this exciting, important vision. It bumped slowly along. With hindsight, a foundation was laid by work in this period on the psychology of programming, exemplified by Weinberg’s influential 1971 book, *The Psychology of Computer Programming*.

Programmers, the principal hands-on users of computers at the time, embodied what would become a key quality: discretionary use. Most programmers used computers because they enjoyed it, unlike data entry and other computer operators. CHI’s focus on discretionary use would differentiate it from HCI research in human factors and ergonomics and management information systems. The latter focused on routinized activity by people for whom computer use was a job requirement. Although a future marked by discretionary use was envisioned early in this period by Engelbart, Ted Nelson, and others, it only became a widespread reality later, with the arrival of home computers and PCs.

ARPA IPTO funding of AI had broadened to include demonstration projects in robotics, speech understanding, and object recognition. In the words of Terry Winograd, whose landmark 1970 AI thesis was for a mathematics degree because of the legacy link of AI to mathematics, these new efforts added an “engineering flavor” to AI research. User

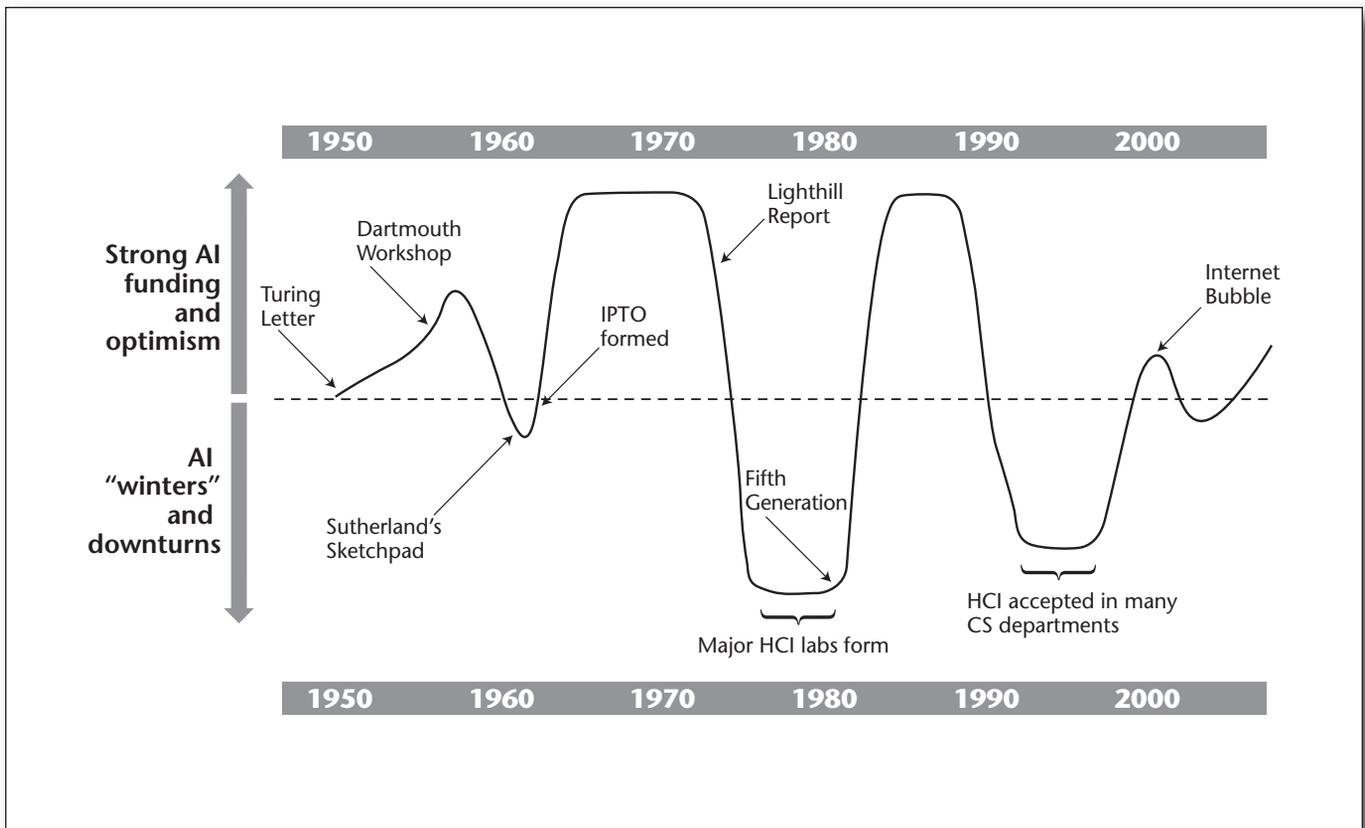


Figure 1. The Changing Seasons of AI and HCI.

Funding climate and public perception with three HCI high points.

studies were undertaken primarily on behalf of trained personnel who would be heavy users of these costly systems, not to appeal to discretionary customers. Such studies were the province of human factors and ergonomics (HFandE), not CHI, so the engineering flavor contributed to the gap between AI and CHI.

The Mid-1970s to the Early 1980s: AI Winter, Springtime for HCI

Eventually, it was clear that AI had been oversold. The bloom was again off the rose, which was to prove to be a hardy perennial, but this time a long downturn resulted. The 1973 British government's Lighthill Report on the prospects of AI was perceived as being so negative that UK funding all but vanished. In the United States, bold ARPA initiatives in timesharing and networking had paid off handsomely. AI ventures had not and were scaled back sharply. The large speech-understanding program was terminated in 1975.

In 1977, *Artificial Intelligence* published an empirical study of natural language understanding by a group of prominent AI and incipient HCI researchers that included Terry Winograd, Donald

Norman, and Danny Bobrow. Initiated prior to the onset of the AI winter, the elegant study identified usability issues and the magnitude of the challenge. With the subsequent dissolution of the collaboration, this may have contributed to the chill.

The AI winter lasted almost a decade, during which human-computer interaction blossomed. Influential HCI research labs formed at PARC, IBM, Digital, the U.K. Medical Research Council Applied Psychology Unit, Bell Labs, and the University of California San Diego. These were central to the growth of SIGCHI, which formed in 1982. The Human Factors Society's Computer Systems Technical Group formed and thrived. HCI also progressed in management information systems.

The UCSD cognitive psychology group contributed particularly heavily. Don Norman, who published HCI-related work in AI forums in the 1970s and 1980s, led university participation in early CHI conferences. UCSD was also central to the early 1980s resurgence of work on "neural nets" or parallel distributed processing models, and hosted the first Cognitive Science conference in 1979. Drawing from computer science, AI, psychology, and other disciplines, cognitive science could be considered a sibling to HCI. Cognitive sci-

ence engaged with HCI almost exclusively only at UCSD and through researchers who spent time there.

By the end of the AI winter, HCI was flourishing. It was not coincidental that HCI waxed as AI waned—or that when fortunes again shifted direction, it was to some extent in tandem.

The 1980s: AI Again Rides High

In 1981 as in 1958, a foreign threat reinvigorated AI funding. Japan, brimming with success in manufacturing, launched the “Fifth Generation” AI effort, built upon the logic-based Prolog language. Reaction was swift. The U.S. Congress amended antitrust laws to permit the founding of the AI-based consortium MCC in 1982. In 1983, U.S. government funding for speech understanding and other AI technologies ramped up sharply, with DARPA again leading the charge. Annual funding for its Strategic Computing Initiative (SCI) rose over time, reaching \$400 million (in today’s U.S. dollars) in 1988. The European ESPRIT and UK Alvey programs began in 1984 with annual AI expenditures of over \$200 million.

New bottles were found, even if some of the wine was old vintage. “AI” was used sparingly; favored were “intelligent knowledge-based systems,” “expert systems,” “knowledge engineering,” “office automation,” “multiagent systems,” and “machine learning.” General problem solving was emphasized less, domain-specific problem solving more. SCI researchers promised to deliver three things: an autonomous land vehicle, a pilot’s associate, and a battlefield management system.

Other factors contributed to the AI boom. In 1982, Symbolics and LMI marketed “LISP machines” optimized for the western AI language of choice. Production systems reached maturity at Carnegie Mellon University and were exported. Neural nets attracted research and media attention.

Once again, some of Turing’s descendants echoed his 1949 claim: machines would soon rival human intelligence, then educate themselves 24/7 and leave *homo sapiens* in the dust. At MCC, Douglas Lenat’s CYC project began in 1984 promising to achieve machine self-education within 10 years. (Some HCI projects proposed at MCC were turned down on the grounds that they did not promise an order of magnitude increase in productivity.) Some neural net proponents promised systems that “could take a bullet and keep going” (despite the fact that the models running ran primarily on single-processor machines). Knowledge engineering would propel businesses forward. In the United States, DARPA, the National Science Foundation, and intelligence agencies were joined by industry in pouring resources into speech recognition and natural language understanding. Annual AI invest-

ment by 150 U.S. companies was an estimated \$2 billion in 1985.

HCI during the AI Summer

SIGCHI formed in 1982 and held its first conference in 1983, inspired by similar conferences held in 1981 and 1982. Despite being cosponsored by ACM, few computer scientists were present; cognitive psychologists and human factors engineers predominated. Then, in 1985, a graphical user interface finally succeeded in the marketplace. The Macintosh computer’s success brought computer science into CHI. GUI elements had been described and prototyped on expensive machines decades earlier by Ivan Sutherland, Douglas Engelbart, and others, but now widely available computers—the focus of CHI—could devote sufficient processing to interaction to greatly expand the design space. Leading computer graphics researchers who were more interested in augmenting interaction than in achieving photorealism shifted from SIGGRAPH to CHI.

CHI also drew in AI researchers whose primary interest had been augmenting human tool use, such as Jim Hollan (then at the University of California San Diego and the Office of Naval Research), whose early work on modeling and visualization was published in AI venues, and Gerhard Fischer, who published work on coaching and critiquing systems in CHI while continuing to publish in AI venues.

As AI gathered momentum in the 1980s, AI researchers and the popular press regularly proclaimed that speech and language understanding would soon be the principal means of communicating with computers. The implication was unmistakable: HCI research focused on other interaction modalities was largely pointless. This was not conducive to disciplinary congeniality.

Two visions competed for government funding, industry investment, and student attention. Research into perceptual-motor and cognitive challenges of GUIs on the one side, exotic machines and glamorous promises of AI on the other. GUIs were cool, but AI dominated funding and media attention and prospered. MCC hired many HCI researchers, including me, to work on LISP projects that included AI-oriented research.

Consider the *International Journal of Man-Machine Studies*. In 1985 its editor, Brian Gaines, published a history of HCI. It ended with a strong focus on the response to the Japanese fifth-generation AI effort. The journal subsequently focused on AI, to the extent it spun off the journal *Knowledge Acquisition* in 1989 (reabsorbing it back into the rechristened *International Journal of Human Computer Studies* in 1996, well into the next AI winter).

CHI never fully embraced speech recognition or natural language, but in the late 1980s papers on

modeling, adaptive interfaces, and advising systems were more numerous. This research required expensive workstations and was carried out in laboratories that equipped everyone with such systems, notably Xerox PARC and MCC.

Governments had an interest in usable AI. They were not only major funders of AI research, they were the major customer of companies that sprang up to market AI products. Use of such products was rarely discretionary. Few people chose to use AI technologies such as AI Corp's Intellect, a natural language interface to databases, or speech and gesture systems delivered to the military. In largely unsuccessful efforts to increase use, the U.S. and UK governments funded considerable work on the human factors of speech systems, expert systems, and knowledge engineering. A typical human factors method involved training followed by controlled experiments and statistical analysis to identify small but significant effects.

In contrast, computer science-oriented CHI researchers, focused on designing for first-time discretionary use, typically used small-group usability tests to identify major effects. Real-world usability was not their focus, which was fortunate given that the technologies were not ready.

The 1990s: Another AI Winter and Period of HCI Growth

As the 1980s ended, DARPA was unhappy. It had spent \$2 billion (in today's dollars) and there were no autonomous land vehicles. There were no pilot's associates or battlefield management systems. Natural language-understanding products foundered in the marketplace. Funding was cut and a long winter set in.

Moravec wrote, "By 1990, entire careers had passed in the frozen winter." For him, the well-funded 1980s were an extension of an intellectual winter due to what he considered to be misguided efforts to use workstations as platforms—not powerful enough to support significant AI. But by most measures, the 1980s had been great years for AI. Attendance at AAAI peaked at 4000 to 5000 from 1986 to 1988. In 1990 it fell below 3000, in 1991 below 2000, and then declined gradually to its current level, slightly under 1000.

There were enough tenured AI professors to ensure that research would survive a downturn, but they scrambled for funding when DARPA pulled out. Much of NSF's Human Interaction Program budget was redirected to AI topics—a program manager told me in the early 1990s that his proudest accomplishment was to double doubling the already sizable funding for natural language understanding. Distributed AI ("multiagent systems") researchers who had been dropped by DARPA quietly convinced a congenial NSF program manager

to redirect a significant part of the collaboration systems budget to them, to the dismay of the associate director when it was discovered. Nevertheless, support for AI graduate students and job-seekers dropped sharply through the 1990s.

In the late 1980s, LISP machines had ceased to provide a significant advantage over LISP running on less expensive platforms, notably the Mac. The collapse of LISP machine companies and PARC's exit from that business had a symbolic impact. The use of LISP itself in AI research declined dramatically through the 1990s. Those in government and industry who had invested in the AI summer were left with very expensive doorstops as reminders.

The drop in support for AI coincided again with a period of HCI growth. Many computer science departments added HCI to their core curricula and hired HCI faculty. The ranks of HCI graduate students swelled as those in AI fell. Fewer AI papers appeared in CHI—one factor was MCC disbanding its large AI-oriented HCI program—but in the 1990s, some new threads of research spanned the fields, utilizing AI techniques on UNIX workstations and personal computers whose prices steadily declined.

Henry Lieberman, whose work on programming by example and interface agents had appeared in AI venues in the 1980s, began publishing in CHI in the 1990s. Recommender system researchers Patti Maes at MIT, Paul Resnick at the University of Michigan, and Joe Konstan, John Riedl, and Loren Terveen at the University of Minnesota published in both venues, some starting in AI and spreading to CHI, others taking the opposite path or publishing in both from the outset. Other CHI publications by researchers with an AI focus were Sharon Oviatt's insightful speech recognition research and applications of machine learning by Eric Horvitz.

However, not many researchers spanned the fields. The first Intelligent User Interface conference in 1993 had considerable participation from human factors and ergonomics, a consequence of the government investment in this approach. When IUI resumed in 1997, it shifted away from human factors and ergonomics/HFandE, toward AI and participation from Europe and Asia, but not toward CHI despite its cosponsorship.

The Current Decade: Supralinear Growth Finally Brings Common Ground

The past 10 years have seen notable AI achievements. A turning point may have been Deep Blue's 1997 defeat of world chess champion Gary Kasparov. Although the victory was bolstered by raw computing power, successful navigation of the middle game required chess analysis of a quality

that surprised many, including long-time AI critic Hubert Dreyfus, author of "What Computers Can't Do."

As the 2000s progressed, remotely controlled robots explored forbidding terrestrial and extraterrestrial terrain, such as the Mars exploration rovers launched in 2003. In 2005, a few vehicles finally made it to the finish line in the DARPA Grand Challenge autonomous land vehicle race. Closer to HCI, machine learning progressed steadily, much of it appearing in IUI. When the rising tide of the Internet floated all boats in the late 1990s, AI contributions such as recommender systems appeared in products. Some were swept away when the tide retreated, but others survived.

As noted above, CHI focuses mainly on technologies when they are first coming into widespread use. For example, GUIs were around from the early 1960s, but only after they succeeded commercially in 1985 did CHI turn its attention to them. Today, supralinear increases in semiconductor capability have reduced the costs of storage, processing, and access to the point where machine learning and other AI technologies can be exploited on standard PCs and through browser access to servers. This all but guarantees that AI will continue to move into the CHI mainstream.

With the Internet boom came another wave of ultraintelligence forecasts by AI luminaries. In 1999, Ray Kurzweil published *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*, and Hans Moravec published *Robot: Mere Machine to Transcendent Mind*. Unlike the 1970 predictions of 5 to 15 years, they estimated 30 to 40. Public anxiety is less than during the Cold War, but we could use a hand with climate change, species extinction, oil dependency, the influence of money on politics, and mortality. Mainstream media find utopian and dystopian forecasts irresistible; for example, a 2009 AAAI study found almost unanimous agreement that ultraintelligence is not on the horizon, but the media nonetheless hyped "the singularity."

Prominent AI researchers at times claim that exaggerated forecasts serve a purpose by attracting and motivating students and funding. Perhaps, but there is a price. Researchers in related fields may consider it intellectually inelegant and unfair competition for resources. Sponsors may eventually react negatively and withdraw support. And it contributes to the "disappearing AI" phenomenon discussed by Lieberman in this issue: the long-standing complaint that once an AI technique succeeds, it is no longer considered AI. However impressive AI achievements are, they are insignificant when contrasted with promises of genius computers that will solve the world's problems and deliver immortality while conversing with us in iambic pentameter.

A Last Challenge: Finding a Tent

If hyperbolic claims can be isolated from the research sphere, the prospects for usable AI and synergy between the fields are outstanding. The last hurdle may be finding a place to meet.

Conferences once focused on community-building, but within U.S. computer science, highly selective conferences that showcase the most polished (or perhaps least controversial) papers have replaced journals as the final repositories of most research. Practitioner attendance has declined. The sense of community suffers when rejection rates are high. Attendance at AAAI, IUI, CHI, and associated conferences peaked years ago. Submission levels have oscillated or plateaued. SIGCHI and SIGART memberships have declined, despite substantial increases in relevant faculty, students, and commercial activity. Conferences are increasingly populated by students, who bring wonderful energy but uncertain commitment.

Many prominent HCI researchers have moved from computer science departments to information schools, introducing new identity issues and distancing them from AI colleagues. AI research is also diffusing into other fields, including but not limited to HCI.

Despite these questions of who we are and where we will meet, technology trends point in the right direction. As costs decline, machine learning can contribute to interaction design and execution, often by focusing on specific user behaviors. Successful AI applications strengthen the tie to HCI by providing research foci as well as by creating a demand for new and better interfaces. AI researchers are acquiring basic HCI skills and more HCI researchers employ AI techniques. Identification of shared purposes, and greater mutual understanding, two goals of this special issue, are indispensable for the next generation of researchers and system builders in both fields.

Acknowledgment

Thanks to the editors of this special issue for encouraging this contribution, and to Don Norman, Anthony Jameson, Brian Butler, and Jim Hollan for comments on early drafts.

Further Reading

Work with a Historical Context

Listed by Author

Gaines, B. R., and Shaw, M. L. G. 1985. From Timesharing to the Sixth Generation: The Development of Human-Computer Interaction, Part I. *International Journal of Man-Machine Studies* 24(1): 1-27.

Grudin, J. 2005. Three Faces of Human-Computer Interaction. *IEEE Annals of the History of Computing* 27(4): 46-62.

Grudin, J. 2006. Turing Maturing: The Separation of Artificial Intelligence and Human-Computer Interaction. *ACM Interactions* 13(5): 54–57.

Grudin, J. 2008. A Moving Target: The Evolution of Human-Computer Interaction. In *Handbook of Human-Computer Interaction*, ed. A. Sears and J. Jacko. Boca Raton, FL: CRC Press.

McCarthy, J. 1988. Book review of *The Question of Artificial Intelligence: Philosophical and Sociological Perspectives* by B. P. Bloomfield. *Annals of the History of Computing* 10(3): 224–229.

Moravec, H. 1998. When Will Computer Hardware Match the Human Brain? *Journal of Evolution and Technology* 1(1).

Norberg, A. L., and O'Neill, J. E. 1996. *Transforming Computer Technology: Information Processing for the Pentagon, 1962–1986*. Baltimore: Johns Hopkins.

Roland, A., and Shiman, P. 2002. *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983–1993*. Cambridge, MA: The MIT Press.

Weinberg, G. 1971 *The Psychology of Computer Programming*. New York: John Wiley & Sons.

Sources on Ultrainelligence

Listed by Author

Asimov, I. 1950. *I, Robot*. New York: Gnome Press.

Darrach, B. 1970. Meet Shaky: The First Electronic Person. *Life Magazine*, November 20.

This article is the source for the Minsky and Quillian quotations.

Dreyfus, H. 1972. *What Computers Can't Do: A Critique of Artificial Reason*. New York: Harper and Row.

Dreyfus identified exaggerated predictions of AI in this book and subsequent revisions, but in dismissing computer chess potential his crystal ball proved no better than that of Simon and Newell, who predicted in 1957 that a computer would be world chess champion by 1967.

Good, I. J. 1965. Speculations Concerning the First Ultrainelligent Machine. *Advances in Computers* 6: 31–88.

Kurzweil, R. 1999. *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. New York: Viking.

Johnson, T. 1985. *Natural Language Computing: The Commercial Applications*. London: Ovum Ltd.

This book is a source for AI investment levels in the mid-1980s.

Lenat, D. 1989. When Will Machines Learn? *Machine Learning* 4: 255–257

Licklider, J. C. R. 1960. Man-Computer Symbiosis. *IRE Transactions of Human Factors in Electronics* HFE-1(1): 4–11.

Markoff, J. 2009. Scientists Worry Machines May Outsmart Man. *New York Times*, July 25, p. A1.

Moravec, H. 1999. *Robot: Mere Machine to Transcendent Mind*. Oxford, UK: Oxford University Press.

Negroponte, N. 1970. *The Architecture Machine: Towards a More Humane Environment*. Cambridge, MA: The MIT Press.

Shannon, C. 1950. Programming a Computer for Playing Chess. *Philosophical Magazine*, Series 7, 41, 314: 256–75.

Simon, H. A. 1960. *The New Science of Management Decision*. New York: Harper.

Simon, H. A., and Newell, A. 1958. Heuristic Problem Solving: The Next Advance in Operations Research. *Operations Research* 6(1): 1–10.

This article contains the text of Simon's address to the 1957 meeting of the Operations Research Society of America.

Turing, A. 1949–1956. Turing's letter to the *Times* was published June 11, 1949. In 1950 he published "Computing Machinery and Intelligence," *Mind* 49: 433–460, which drew more attention when republished in 1956 as "Can a Machine Think?" In J. R. Newman, ed., *The World of Mathematics*, volume 4, 2099–2123. New York: Simon and Schuster.

Sources Bridging the Gap between AI and CHI

Listed by Date of Publication

These researchers crossed or spanned the disciplines, indicating that divisive forces can be overcome.

The 1970s

Thompson, H.; Winograd, T.; Bobrow, D.; Norman, D.; Kay, M.; and Kaplan, R. 1977. GUS, A Frame-Driven Dialog System. *Artificial Intelligence* 8(2): 155–173.

Well worth reading today.

Norman, D. A. 1978. Analysis and Design of Intelligent Systems. In *Human and Artificial Intelligence*, ed. F. Klix. Berlin, DDR: Veb Deutscher Verlag der Wissenschaften.

The 1980s

Hollan, J. D.; Hutchins, E. L.; and Weitzman, L. 1984. Steamer: An Interactive Inspectable Simulation-Based Training System. *AI Magazine* 5(2): 15–27.

These researchers eventually focused on HCI, as in the influential chapter cited below.

Lieberman, H. 1984. Seeing What Your Programs Are Doing. *IJMMMS* 21(4): 311–331.

Through 1989, Lieberman published in graphics and AI venues, but with user-centered work such as this. Starting in 1989 with a tutorial on user interfaces for program debugging Lieberman became a prolific contributor to CHI, while still publishing in graphics and AI literatures.

Fischer, G.; Lemke, A.; and Schwab, T. 1985. Knowledge-Based Help Systems. In *Human Factors in Computing Systems: Proceedings of CHI'85*, 161–167. New York: Association for Computing Machinery.

Hutchins, E. L.; Hollan, J. D.; and Norman, D. A. 1986. Direct Manipulation Interfaces. In *User Centered System Design*, ed. D. A. Norman and S. W. Draper, 87–124. Mahwah, NJ: Erlbaum.

Fischer, G. 1987. A Critic for LISP. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 177–184. San Francisco: Morgan Kaufmann.

"Critic" refers to software that inspects a program and suggests improvements. Fischer has continued to contribute to HCI and AI (as well as software engineering).

In the late 1980s, the AI boom drew in HCI researchers. An ambitious MCC effort headed by Jim Hollan produced works such as the following:

Terveen, L. G., and Wroblewski, D. A. 1990. A Collaborative Interface for Browsing and Editing Large Knowledge Bases. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 491–496. Menlo Park, CA: AAAI Press.

Hollan, J. D.; Rich, E.; Hill, W. C.; Wroblewski, D.; Wilner, W.; Wittenburg, K.; Grudin, J.; and the MCC Human Interface Lab Members. 1991. An Introduction to HITS: Human Interface Tool Suite. In *Intelligent User Interfaces*, ed. J. Sullivan, 293–338. Reading, MA: Addison Wesley.

The 1990s

Shardanand, U., and Maes, P. 1995. Social Information Filtering: Algorithms for Automating “Word of Mouth.” In *Human Factors in Computing Systems: Proceedings of CHI’95*, 210–217. New York: Association for Computing Machinery.

Maes works primarily in AI but was influential in bringing social filtering research to the attention of the HCI community.

Norman, D. A. 1997. How Might People Interact with Agents. In *Software Agents*, ed. J. M. Bradshaw. Menlo Park, CA: AAAI Press.

This is a modified version of Norman’s 1994 article in *Communications of the ACM* 37: 68–71.

Resnick, P. 1997. Filtering Information on the Internet. *Scientific American*, March, 106–108.

Resnick had previously published in CHI and AI conferences, including work related to this at SIGHCI-sponsored CSCW’94.

Good, N.; Schafer, J. B.; Konstan, J. A.; Borchers, A.; Sarwar, B. M.; Herlocker, J. L.; and Riedl, J. 1999. Combining Collaborative Filtering with Personal Agents for Better Recommendations. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, 439–446. Menlo Park, CA: AAAI Press.

This primarily Minnesota-based group published work on collaborative filtering extensively at SIGCHI-sponsored conferences in the 1990s. A more recent example follows.

Horvitz, E. 1999. Principles of Mixed-Initiative User Interfaces. In *Human Factors in Computing System: Proceedings of CHI’99*, 159–166. New York: Association for Computing Machinery.

Another leading figure in AI (and a recent president of AAAI), Horvitz is engaged with the HCI literature:

Oviatt, S. 1999. Mutual Disambiguation of Recognition Errors in a Multimodal Architecture. In *Human Factors in Computing System: Proceedings of CHI’99*, 576–583. New York: Association for Computing Machinery.

Oviatt works primarily in AI, but has published elegant work applying HCI approaches to AI systems in novel ways.

The Current Decade

Although work in the current decade is too extensive to list, an example alluded to in the article is the following:

Herlocker, J. L.; Konstan, J. A.; Terveen, L. G.; and Riedl, J. T. 2004. Evaluating Collaborative Filtering Recommender Systems. *ACM Transactions on Information Systems* 22(1) 5–53.

AAAI E-mail Addresses

Please note that AAAI will be modifying its e-mail addresses in 2010 in an effort to reduce the amount of spam that we are receiving. We will be adding 10 to all e-mail addresses, as follows:

aaai10
 aiide10
 aimagazine10
 aitopics10
 fellows10
 fss10
 iaai10
 icwsm10
 info10
 membership10
 orders10
 press10
 sss10
 volunteer10
 workshops10

The number will be updated on an annual basis. AAAI can also be reached by filling out the contact form at www.aaai.org/scripts/Contact/contact.php.

Jonathan Grudin has a BA in mathematics-physics from Reed College and a Ph.D. in cognitive psychology from the University of California San Diego. Prior to joining Microsoft Research, where he is a principal researcher, he was a professor of information and computer science at the University of California Irvine. He is the associate editor for human-computer interaction for *ACM Computing Surveys* and writes and edits a column on history topics for *ACM Interactions* magazine.