

Stanford University Program in
Human-Computer Interaction

From Computing Machinery to Interaction Design

Terry Winograd
Stanford University

Published online by permission from Peter Denning and Robert Metcalfe (eds.), *Beyond Calculation: The Next Fifty Years of Computing*, Springer-Verlag, 1997, 149-162.

When asked to project 50 years ahead, a scientist is in a bit of a quandary. It is easy to indulge in wishful thinking, or to promote favorite current projects and proposals, but it is a daunting task to anticipate what will actually come to pass in a time span that is eons long in our modern accelerated age. If fifty years ago, when the ACM was founded, biologists had been asked to predict the next 50 years of biology, it would have taken amazing prescience to anticipate the science of molecular biology. Or for that matter, only a few years before the initiation of the ACM even those with the most insight about computing would have been completely unable to foresee today's world of pervasive workstations, mobile communicators, and gigabit networking.

1. Directions of change

In order to get some grip on the future we need to take advantage of the past and the present: to identify some trajectories that are already in swing, and to look at some further points towards which they are headed. The method isn't foolproof---we can always be surprised. But at least it gives us some initial starting points. So we will begin by examining three trends whose directions are already visible, and then project an image of where they may lead in the coming half-century. For convenience, let us label these trajectories:

- Computation to communication
- Machinery to habitat
- Aliens to agents

1.1. From computation to communication

When digital computers first appeared a half-century ago, they were straightforwardly viewed as "machinery for computing." A computer could make short work of a task such as calculating ballistics trajectories or breaking codes, which previously required huge quantities of computation done by teams of human "computers." Even a quarter century later, when the Internet was created, the network was seen primarily as a tool for facilitating remote computation (for a history of the Internet, see the [article by Hafner](#)). On the net, a computing task could be carried out using a computer that was physically removed from the person who needed the work done and who controlled its activity.

With the recent -- and quite sudden -- emergence of mass-appeal Internet-centered applications, it has become glaringly obvious that the computer is not a machine whose main purpose is to get a computing task done. The computer, with its attendant peripherals and networks, is a machine that provides new ways for people to communicate with other people. The excitement that infuses computing today comes from the exploration of new capacities to manipulate and communicate all kinds of information in all kinds of media, reaching new audiences in ways that would have been unthinkable before the computer.

In retrospect, it didn't take the World Wide Web to show us the centrality of communication. There are many earlier points in the development of computers that dramatically revealed the precedence of communication over computation. One obvious clue could be seen in the pattern of adoption of the Internet itself. Instead of using the net for remote computing, as in the original proposals, the vast majority of people with Internet connections have used it to communicate with each other: through email, through newsgroups, through real-time "talk," for example. Whole new families of communication-centered applications have appeared, from "groupware" to MUDs and MOOs.

The story is the same for the "personal" computer. The suites of applications that dominate the office market today consist primarily of communication tools: word processors, presentation programs, email, file sharing, and contact managers. Even the one apparent exception -- the spreadsheet -- is used more prominently for communicating results than for calculating them.

In some sense this should be no surprise, given what we can observe about human nature. People are primarily interested in other people, and are highly motivated to interact with them in whatever media are available.

New technologies, from the telegraph to the World Wide Web have expanded our abilities to communicate widely, flexibly, and efficiently. This urge to communicate will continue to drive the expanding technology with the advent of widespread two-way video, wireless connectivity, and high-bandwidth audio, video, 3-D imaging, and more yet to be imagined.

Within the computing industry we are now seeing a new emphasis on communication, reflected in a concern for "content." The companies that can put extensive resources into computing-system development are shifting their gaze away from what the machine "does" towards what it "communicates". As an illustrative series of points along our trajectory, consider the history of Microsoft, which began with operating systems, then expanded into the world of software applications, and now is actively moving into the content arena with a joint TV effort with NBC, an on-line magazine, and major acquisitions of visual materials. The latter may not be under the Microsoft label, but new companies by the founders Bill Gates (Corbis, a graphical image library) and Paul Allen (Starwave, an on-line information and entertainment service), give a sense of where the big money is heading. This shift towards content extends beyond the software companies, such as Microsoft, reaching to those that have been primarily hardware companies. Workstation maker Silicon Graphics is moving into the entertainment business, and chip maker Intel has recently opened a new research laboratory with interest in high-level questions about human communication and uses of computers in the home.

There will always be a need for machinery and a need for software that runs on the machinery, but as the industry matures, these dimensions will take on the character of commodities, while the industry-creating innovations will be in what the hardware and software allow us to communicate.

1.2. From machinery to habitat

In the early days of computing, the focus of computer scientists was---as the original name of the ACM implies---on the "machinery." In order to use a computer, one needed to understand how to give it instructions. As the field matured, the details of the physical machinery were gradually submerged beneath the surface of the software: the higher-level expressions of the computation to be performed, separated from the details of how an electronic device would operate to carry out such computations. Most people who use computers today (with the obvious exception of computer scientists) do not know or care much about the details of the processor, its architecture, or the ways in which it operates. If you ask people what computer they use, they will often say "Windows." Going one step further, many people will say that the computer they use is "Microsoft Word" or "Netscape" without even distinguishing among the operating-system platforms on which the software is executing.

It is easy for experts to sneer at this kind of technical inaccuracy as a symptom of ignorance or misunderstanding. But in fact, it is quite logical and appropriate for many kinds of computer users. I could ask a relatively sophisticated computer user whether he or she has a computer with NMOS or CMOS transistors and draw a complete blank. For the chip manufacturer, this is an important distinction, but for the average user (even the average computer professional), the distinction only matters because it is reflected in performance parameters: How fast is the machine? How much electric power does it require? The domain in which NMOS and CMOS are distinguished isn't the domain in which the computer user experiences its activity. In just that same manner, whether a machine is PowerPC or Wintel, whether it runs Windows95 or MacOS, isn't in the domain in which the average user is coming to experience the computer. Differences at this level can have impact in terms of speed, cost, the availability of new software, and the like, but they are only of relevance in this indirect way.

With the Web we are seeing the distancing from the machine taken yet a step further. In spite of tremendous efforts by Netscape and Microsoft to differentiate their browsers, users will end up caring only indirectly about what software is running on their machines at all. Their experience isn't of a machine, or a program, but of entering into the reaches of a cyberspace populated with text, graphics, video, and animations, or even more to the point, consisting of catalogs, advertisements, animated stories, and personal picture galleries.

The word "cyberspace" is often bandied about as a symbol of the new computing, and it has become a trendy cliché. But it has a more significant meaning than that. The fact that cyberspace is termed a "space" reflects a deep metaphor, of the kind that [Lakoff and Johnson](#) say we "live by". A space is not just a set of objects or activities, but is a medium in which a person experiences, acts, and lives. The ultimate extension of cyberspace is depicted in "cyberpunk" science fiction works such as William Gibson's [Neuromancer](#) and Neal Stephenson's [Snowcrash](#). In these bleak visions of the future, technologies of virtual reality enable the protagonists to live in a virtual world that is in many ways more real than the physical world from which they are "jacked in." For the voyager in this cyberspace there are no computers, no software in sight---just the "metaverse," which is experienced through the operation of necessary but ultimately invisible computing machinery.

In spite of all the popular excitement about virtual reality, an immersive high-fidelity 3-D display with gloves and goggles isn't a necessary (or even a major) component of creating a captivating virtual space. Over the past decade a menagerie of virtual worlds have evolved which offer their

denizens nothing more than plain text, typed and read on a standard display. In these MUDs, MOOs, and IRC (Internet Relay Chat), people carry on their professional, recreational, and private lives in a way that challenges our basic notions about what is "real." As one of her informants said to Sherry Turkle, author of [Life on the Screen](#), "real life is just one more window."

Turkle's informant may be extreme, as are some of the others she interviewed who spend most of their waking hours online, meeting and falling in love with people, having fights and sex and generally doing all of the things that make up our everyday lives. But anyone who uses computers extensively is aware of the feeling that when we are online we are located in cyberspace as surely as we are located in physical space. When I carry my laptop computer with me on a trip and connect to the network from a hotel room, I feel very much that I am at my office. But if I am physically in my office but the computer is down, I feel that I am not at my normal workplace. In a very real sense, the "webhead" who spends the night surfing from site to site has been traveling in cyberspace, not sitting in a chair in her dorm room or computer laboratory.

The traditional idea of "interface" implies that we are focusing on two entities, the person and the machine, and on the space that lies between them. But beyond the interface, we operate in an "interspace" that is inhabited by multiple people, workstations, servers, and other devices in a complex web of interactions. In designing new systems and applications, we are not simply providing better tools for working with objects in a previously existing world. We are creating new worlds. Computer systems and software are becoming media for the creation of virtualities: the worlds in which users of the software perceive, act, and respond to experiences.

1.3. Aliens to agents

The cyberpunk novels that glorify life in cyberspace offer an interesting contrast to an older and still prevalent vision of the computer future, one in which computing machines come to be independent, alien beings in the world. From ancient tales of the golem to Arthur Clarke's [2001](#) and Isaac [Asimov's robots](#), people have been fascinated by the prospect of sharing our physical and mental worlds with a new species of beings that we ourselves have created. The field of artificial intelligence (AI) was born from this vision, and its founding leaders made bold predictions of "ultra-intelligent machines" and the inevitability of building computers that would soon duplicate and then surpass the human intellect.

From the vantage point of AI researchers of twenty-five years ago, it would have been ludicrous to predict that as we approached the year 2001, almost all of the companies devoted to AI technologies would have faded from existence, while the hot new mega-company that shattered Wall Street records was Netscape, a producer of better network terminals (as browsers would have been called in those days). And right behind it were companies such as Yahoo, whose only product is an index of things to be found on the net. What happened to the intelligent computers?

In the early days of artificial intelligence, researchers' vision was focused on quasi-intelligent beings that would duplicate a broad range of human mental capacities. Even such apparently simple mental feats such as walking, seeing what is around us, and employing common sense turned out to be surprisingly hard. As the field matured into the engineering of "expert systems" in the 1980s, the focus shifted from "intelligence" to "knowledge." The large investments and research projects moved to creating large knowledge bases that would encode human expert knowledge for computers to serve as skilled, narrow, and presumably submissive experts in practical areas of commerce and science. When the promises of massively increased productivity

through [knowledge engineering](#) didn't come true, expert systems joined ultra-intelligent machines in the category of "not practical with today's technology." Expectations gradually became more modest, concentrating on the development of "intelligent agents," which apply AI technologies in limited ways to help people interact with computer systems.

Today's popular press plays up efforts like those of [Pattie Maes](#) and her research group at the MIT Media Laboratory, where they have produced agents to help people browse the web, choose music, and filter email. In fact, a notable indicator of the current trajectory is the ascendancy of MIT's Media Lab, with its explicit focus on media and communication, over the AI Laboratory, which in earlier days was MIT's headline computing organization, one of the world centers of the original AI research.

With hindsight, of course, it is easy to fault early predictions and quixotic enterprises, such as [Lenat's attempt](#) to produce common sense in computers by encoding millions of mundane facts in a quasi-logical formalism. But we can sympathize with the optimistic naivete of those whose predictions of future computing abilities were based on projecting the jump that led us from almost nothing to striking demonstrations of artificial intelligence in the first twenty-five years of computing. A straightforward projection of the rate of advance seemed that it would lead within another few decades to fully intelligent machines.

But there is something more to be learned here than the general lesson that curves don't always continue going up exponentially (a lesson that the computing field in general has yet to grapple with). The problem with artificial intelligence wasn't that we reached a plateau in our ability to perform millions of LIPS (logical inferences per second), or to invent new algorithms. The problem was in the foundations on which the people in the field conceived of intelligence.

As was pointed out from early in the history of computers by philosophers such as [Hubert Dreyfus](#), the mainstream AI effort rested on a view of human intelligence and action that implicitly assumed that all of intelligence could be produced by mechanisms that were inherently like the conscious logical manipulation of facts represented in language. This chapter is not the appropriate place to present the detailed arguments (see Winograd and Flores, [Understanding Computers and Cognition](#) for an extended discussion). What is relevant to our analysis here is that what appeared to be inevitable trends were based on misconceptions about intellectual foundations. Although philosophically based critiques were considered imprecise and irrelevant by most people in computing, the passage of time has made them seem more revealing and predictive than the more obvious technical indicators that were visible at the time.

These critiques focused on the central body of technological development that is sometimes referred to as "Good Old Fashioned AI (GOFAI)". What about neural nets, genetic programming, self-organizing systems, and other nontraditional approaches to producing intelligent computers? Although the practical results from these efforts are as of yet rather meager, their proponents reflect an attitude that is at the center of technological progress. The biggest advances will come not from doing more and bigger and faster of what we are already doing, but from finding new metaphors, new starting points. Of course, most of these will fail and we cannot tell in advance which ones will lead to a surprising success or how long it will be until something good shows up. But there are ways to open up a clearing in which new possibilities can be glimpsed, even if their full potential cannot be known. The message from the history AI is that we need to be prepared to reexamine our foundational assumptions and start from new footings.

2. The ascendancy (and independence) of interaction design

Among the many possibilities suggested by these three trajectories, one seems particularly relevant to the ACM and its role in the next fifty years of computing. In a way the prediction is paradoxical -- the computing field will become larger, and at the same time the computing profession will become narrower.

In the next fifty years, the increasing importance of designing spaces for human communication and interaction will lead to expansion in those aspects of computing that are focused on people, rather than machinery. The methods, skills, and techniques concerning these human aspects are generally foreign to those of mainstream computer science, and it is likely that they will detach (at least partially) from their historical roots to create a new field of "interaction design."

2.1. The shifting boundaries

As the trajectories depicted here continue along their current courses, the computing industry will continue to broaden its boundaries -- from machinery to software to communication to content. The companies that drive innovation will not be those that focus narrowly on technical innovation but those that deal with the larger context in which the technologies are deployed.

As the focus of commercial and practical interest continues to shift, so will the character of the people who will be engaged in the work. Many of the most exciting new research and development in computing will not be in traditional areas of hardware and software but will be aimed at enhancing our ability to understand, analyze, and create interaction spaces. The work will be rooted in disciplines that focus on people and communication, such as psychology, communications, graphic design, and linguistics, as well as in the disciplines that support computing and communications technologies.

As computing becomes broader as a social and commercial enterprise, what will happen to computer science as a professional discipline? Will it extend outward to include graphic design, linguistics, and psychology? What would it even mean to have a science of that breadth?

It is more realistic to imagine that computer science will not expand its boundaries, but will in fact contract them while deepening its roots. Much of the commercial success of computing-related industries will be driven by considerations outside of the technical scope of computer science as we know it today, but there will always be new theories, discoveries, and technological advances in the hardware and software areas that make up the core of the traditional discipline.

As an analogy, consider the role of mechanical engineering and thermodynamic theory in the world of the automobile (or, more broadly, of transportation vehicles). It is clear that success in today's automotive market is determined by many factors that have little to do with science and engineering. They range from the positioning of a vehicle in the market (consider the rise of four-wheel drive sports vehicles) to the ability to associate it with an appealing emotional image through styling, furnishing, and advertising. Engineering is still important and relevant, but it isn't the largest factor for success, and it isn't the dominating force in the automobile industry.

We can expect the same kind of decoupling in the computer world. The flashy and immensely lucrative new startup companies will depend less on new technical developments and more on the kinds of concerns that drive the film industry or the automobile industry. The computing industry will come to encompass work from many different professions, one of which will be the computer

science profession, which will continue to focus on the computing aspects that can be best approached through its formal theories and engineering methods. As the center of action shifts, computer science may lose some of its current highly favorable funding position, but it will gain in its intellectual coherence and depth. To put it simply, it will become easier in the future to see the difference between a significant scientific insight and a hot new product -- a distinction that is blurred in today's world of proliferating technology startups and product-driven research funding.

2.2. The Emergence of Interaction Design

In portraying the broadening scope of computing, I have alluded to many existing disciplines, ranging from linguistics and psychology to graphic and industrial design. Human-computer interaction is by necessity a field with interdisciplinary concerns, since its essence is interaction that includes people and machines, virtual worlds and computer networks, and a diverse array of objects and behaviors.

In the midst of this interdisciplinary collision, we can see the beginnings of a new profession, which might be called "interaction design." While drawing from many of the older disciplines, it has a distinct set of concerns and methods. It draws on elements of graphic design, information design, and concepts of human-computer interaction as a basis for designing interaction with (and habitation within) computer-based systems. Although computers are at the center of interaction design, it is not a subfield of computer science.

As a simple analogy, consider the division of concerns between a civil engineer and an architect as they approach the problem of building a house or an office building. The architect focuses on people and their interactions with and within the space being created. Is it cozy or expansive? Does it provide the kind of spaces that fit the living style of the family or business for whom it is being designed? What is the flow of work within the office, and what kinds of communication paths does it depend on? How will people be led to behave in a space that looks like the one being designed? Will the common areas be ignored, or will they lead to increased informal discussion? What are the key differences between designing a bank and a barbershop, a cathedral and a cafe?

The engineer, on the other hand, is concerned with issues such as structural soundness, construction methods, cost, and durability. The training of architects and engineers is correspondingly different. Architects go through an education in a studio environment that emphasizes the creation and critique of suitable designs. Engineering emphasizes the ability to apply the accumulated formal knowledge of the field to be able to predict and calculate the technical possibilities and resource tradeoffs that go into deciding what can be constructed.

As with a house or an office building, software is not just a device with which the user interacts; it is also the generator of a space in which the user lives. Interaction design is related to software engineering in the same way architecture is related to civil engineering.

Although there is no clear boundary between design and engineering, there is a critical difference in perspective (see Terry Winograd, [Bringing Design to Software](#)). All engineering and design activities call for the management of tradeoffs. In classical engineering disciplines, the tradeoffs can often be quantified: material strength, construction costs, rate of wear, and the like. In design disciplines, the tradeoffs are more difficult to identify and to measure because they rest on human needs, desires, and values. The designer stands with one foot in technology and one foot in the domain of human concerns, and these two worlds are not easily commensurable.

As well as being distinct from engineering, interaction design does not fit into any of the existing design fields. If software were something that the computer user just looked at, rather than operated, traditional visual design would be at the center of software design. If the spaces were actually physical, rather than virtual, then traditional product and architectural design would suffice. But computers have created a new medium -- one that is both active and virtual. Designers in the new medium need to develop principles and practices that are unique to the computer's scope and fluidity of interactivity.

Architecture as we know it can be said to have started when the building technologies, such as stone cutting, made possible a new kind of building. Graphic design emerged as a distinct art when the printing press made possible the mass production of visual materials. Product design grew out of the development, in the 20th century, of physical materials such as plastics, which allowed designers to effectively create a vastly increased variety of forms for consumer objects. In a similar way, the computer has created a new domain of possibilities for creating spaces and interactions with unprecedented flexibility and immediacy. We have begun to explore this domain and to design many intriguing objects and spaces, from video games and word processors to virtual reality simulations of molecules. But we are far from understanding this new field of interaction design.

A striking example at the time of this writing is the chaotic state of "web page design". The very name is misleading, in that it suggests that the World Wide Web is a collection of "pages," and therefore that the relevant expertise is that of the graphic designer or information designer. But the "page" today is often much less like a printed page than a graphic user interface -- not something to look at, but something to interact with. The page designer needs to be a programmer with a mastery of computing techniques and programming languages such as Java. Yet, something more is missing in the gap between people trained in graphic arts and people trained in programming. Neither group is really trained in understanding interaction as a core phenomenon. They know how to build programs and they know how to lay out text and graphics, but there is not yet a professional body of knowledge that underlies the design of effective interactions between people and machines and among people using machines. With the emergence of interaction design in the coming decades, we will provide the foundation for the "page designers" of the future to master the principles and complexities of interaction and interactive spaces.

The mastery we can expect is, of course, incomplete. Taking seriously that the design role is the construction of the "interspace" in which people live, rather than an "interface" with which they interact, the interaction designer needs to take a broader view that includes understanding how people and societies adapt to new technologies. To continue with our automotive analogy, imagine that on the fiftieth anniversary of the "Association for Automotive Machinery" a group of experts had been asked to speculate on the "the next fifty years of driving." They might well have envisioned new kinds of engines, automatic braking, and active suspension systems. But what about interstate freeways, drive-in movies, and the decline of the inner city? These are not exactly changes in "driving," but in the end they are the most significant consequences of automotive technology.

Successful interaction design requires a shift from seeing the machinery to seeing the lives of the people using it. In this human dimension, the relevant factors become hard to quantify, hard to even identify. This difficulty is magnified when we try to look at social consequences. Will the computer lead to a world in which our concept of individual privacy is challenged or changed? Will online addiction become a social problem to rival drug use? Will political power gravitate

towards people or institutions who have the most powerful communications technologies or who aggregate control over media? Will there be a general turning away from computing technologies in a "back-to-nature" movement that reemphasizes our physical embodiment in the world?

All of these and many more futures are possible, and they are not strictly determined by technological choices. There is a complex interplay among technology, individual psychology, and social communication, all mixed in an intricate chaotic system. Details that seem insignificant today may grow into the major causal factors in over the next fifty years. Trends that seem obvious and inevitable may be derailed for what currently appear to be insignificant reasons.

As with the technological predictions, we need to accept the unpredictability of changes in the social dimension without abandoning our attempts to see, and guide, where things are going. Perhaps fifty years is too long a span across which to see clearly, and we can ask what is happening in the next ten or twenty years, or even this year or this month. Many of the concerns that are dimly visible in the future have concrete reflections in today's society. Many of them are more highly visible to those who have an understanding of the theoretical and practical potentials of new computing technologies, and therefore we as professionals with relevant areas of expertise have a special responsibility to point out both possibilities and dangers.

Interaction design in the coming fifty years will have an ideal to follow that combines the concerns and benefits of its many intellectual predecessors. Like the engineering disciplines, it needs to be practical and rigorous. Like the design disciplines, it needs to place human concerns and needs at the center of guiding design; and like the social disciplines, it needs to take a broad view of social possibilities and responsibilities. The challenge is large, as are the benefits. Given the record of how much computing has achieved in the last fifty years, we have every reason to expect this much of the future.

References

1. Asimov, Isaac. *I, Robot*. New York: New American Library Of World Lit., 1950; *The Foundation Trilogy*. Garden City, N. Y.,: Doubleday [1951-53]; *The Complete Robot*. Garden City, N.Y.: Doubleday, 1982; *Robots, Machines In Man's Image*. New York: Harmony Books, 1985.
2. Clarke, Arthur, 2001: *A Space Odyssey*, London: Hutchinson/Star, 1968.
3. Dreyfus, Hubert L. *What Computers Can't Do* (1st ed.). New York, Harper & Row, 1972.
4. Feigenbaum, Edward A. And Pamela McCorduck, *The Fifth Generation*. Reading, MA: Addison-Wesley. 1983.
5. Gibson, William. *Neuromancer* . New York: Ace Science Fiction Books, 1984.
6. Hafner, Katie, and Matthew Lyon, *Where Wizards Stay Up Late: The Origins of the Internet*, New York: Simon and Schuster, 1996.
7. Lakoff, George, and Mark Johnson, *Metaphors We Live By*, Chicago: Univ. of Chicago Press, 1980.

8. Lenat, Douglas B. *Building Large Knowledge-Based Systems*. Reading, MA: Addison-Wesley, 1990.
9. Maes, Pattie, *Designing Autonomous Agents*. Cambridge, MA: MIT Press, 1990.
10. Stephenson, Neal, *Snow Crash*. New York: Bantam, 1993.
11. Turkle, Sherry, *Life on the Screen: Identity in the Age of the Internet*. New York: Simon and Schuster, 1995.
12. Winograd, Terry, with John Bennett, Laura De Young, and Bradley Hartfield (eds.), *Bringing Design to Software*, Reading, MA: Addison-Wesley, 1996.
13. Winograd, Terry and Fernando Flores, *Understanding Computers and Cognition: A New Foundation for Design*. Norwood, NJ: Ablex, 1986. Paperback issued by Addison-Wesley, 1987.
14. Winograd, Terry, "Thinking machines: Can there be? Are We?," in James Sheehan and Morton Sosna, eds., *The Boundaries of Humanity: Humans, Animals, Machines*, Berkeley: University of California Press, 1991, pp. 198-223.