


DÆDALUS

JOURNAL OF THE AMERICAN ACADEMY OF ARTS AND SCIENCES

Artificial Intelligence

Winter 1988

Issued as Volume 117, Number 1, of the *Proceedings
of the American Academy of Arts and Sciences*



Winter 1988, "Artificial Intelligence." Issued as Volume 117, Number 1, of the Proceedings of the American Academy of Arts and Sciences.

© 1988 by the American Academy of Arts and Sciences.

Library of Congress Catalog Number 12-30299.

Editorial Offices: Dædalus, Norton's Woods, 136 Irving Street, Cambridge, MA 02138.

Dædalus (ISSN 0011-5266) is published quarterly by the American Academy of Arts and Sciences. Subscription rate: \$16, one year; \$24, two years; \$32, three years.

All subscription orders, single copy orders, and change of address information must be sent in writing to the Dædalus Business Office, P.O. Box 515, Canton, MA 02021, U.S.A. (Telephone: 617-828-8450). Orders for single copies must be accompanied with prepayment of \$5.00.

Printed in the United States of America.

Printing Office: 2901 Byrdhill Road, Richmond, VA 23228.

U.S.A. newsstand distribution by Eastern News Distributors, Inc., 1130 Cleveland Road, Sandusky, OH 44870.

Second-class postage paid at Canton, Massachusetts, and at additional mailing offices.



Sherry Turkle

Artificial Intelligence and Psychoanalysis: A New Alliance

ARTIFICIAL INTELLIGENCE and psychoanalysis appear to be worlds apart. Psychoanalysis looks for what is most human: the body, sexuality, what follows from being born of a woman and raised in a family. Artificial intelligence looks deliberately for what is least specifically human: the foundation of its theoretical vision is the thesis that the essence of mental life is a set of principles that could be shared by people and machines.¹

There is another way in which they appear worlds apart. Artificial intelligence seems scientifically ascendant and has increasingly determined the agenda for academic psychology through its influence on cognitive science. In contrast, psychoanalysis is rejected by academic psychology and in conflict with dominant biological trends in psychiatry. Although there have been recent flurries of interest in Freudian theory, they have come from the worlds of literary analysis and philosophy. To scientific circles, psychoanalysis appears a frozen discipline—frozen in the scientific language of another time, frozen in the psychological assumptions of another culture.

In this essay I propose that if psychoanalysis is in trouble, artificial intelligence may be able to help. And I suggest the nature of this help by arguing that one of the ways computers influence psychological thinking is through a route that is not essentially technical. Rather, computers provide sciences of mind with a kind of theoretical

Sherry Turkle is associate professor of sociology in the Program in Science, Technology, and Society at the Massachusetts Institute of Technology. She is the author of Psychoanalytic Politics: Freud's French Revolution and of The Second Self: Computers and the Human Spirit.

legitimation that I call sustaining myths. Indeed, the early impact of the computer on psychology was clearly of this nature.

SUSTAINING MYTHS

As recently as the 1950s, behaviorism dominated American academic psychology, its spirit captured by saying that it was permissible to study remembering but considered a violation of scientific rigor to talk about "the memory." One could study behavior but not inner states. The study of mind had to be expressed in terms of stimulus and response. In today's jargon, what lay between was a black box that must not be opened even speculatively.

By the end of the 1960s, the behaviorist hegemony was broken, as were inhibitions about the study of memory and the inner processes of mind. Indeed, within academic psychology scarcely a trace remained of behaviorist methodology. Behaviorism had not been refuted by a critical experiment. There had been many factors influencing this scientific revolution, including the political and cultural climate of the 1960s. And one of the most central was the computer.

The computer's role in the demise of behaviorism was not technical. It was the very *existence* of the computer that provided legitimation for a radically different way of seeing mind. Computer scientists had of necessity developed a vocabulary for talking about what was happening inside their machines, the "internal states" of general systems. If the new machine "minds" had inner states, surely people had them too. The psychologist George Miller, who was at Harvard during the heyday of behaviorism, has described how psychologists began to feel embarrassed about not being allowed to discuss memory now that computers had one:

The engineers showed us how to build a machine that has memory, a machine that has purpose, a machine that plays chess, a machine that can detect signals in the presence of noise, and so on. If they can do that, then the kind of things they say about the machines, a psychologist should be permitted to say about a human being.²

The computer presence relegitimated the study of memory and inner states within scientific psychology. Many technical concepts

that psychologists picked up from computation—ideas from cybernetics and automata theory—had existed before real computers but became more compelling because of them. “Suddenly,” says Miller, “engineers were using the mentalistic terms that soft-hearted psychologists had wanted to use but had been told were unscientific.”³ Computational ideas, computational language, and the physical presence of the machines all created an intellectual climate in which it was permissible to talk about mental processes banned by behaviorism. The computer presence served as a sustaining myth for a new psychology of inner states that came to be known as cognitive science.

Computer programs provided a way to discuss beliefs and rules as causing behavior. Why did pawn take pawn? Earlier psychologies would have rejected “because the pawn blocked the bishop” as a causal explanation. It would merely be giving the chess player’s “reasons.” But if mind is program, reasons become explanations. A large part of the computer’s appeal for psychologists is that it allows them to open the black box that is the mind. Once that box is open, the computer suggests ways to fill it with concepts that are close to commonsense understandings.

Indeed, the crux of Miller’s story about memory is that computers gave psychologists permission to investigate something that “everyone knows” but that had been banished from science—the idea that people have memories. In the past two decades, cognitive science has been dominated by the computer legitimating the study of something else that “everyone knows,” this time the idea that people have information and use rules, and that much of this information can be formulated in words. In the late 1950s Allen Newell and Herbert Simon built a computer program called the General Problem Solver (GPS), which was guided by something very close to verbal reasons recoded as computational rules. Questions such as “Why did GPS do such and such?” could be answered by reference to whatever rules it had been given. Why should references to rules not be used to answer questions about what people do when faced with similar problems? The existence of GPS gave credibility to the question.

There is a widespread view that the computer presence tends to move psychology toward more rigorous and quantifiable theories, arguing that the computer, by its nature, requires rules, rigor, and formalism. But the story of the computer’s influence on psychology is

not so simple. For example, its "first act"—the attack on behaviorism—went in the direction of creating a less rather than a more constrained, a "softer" rather than a "harder" science of mind.

Artificial intelligence is the most explicit channel for the computer's influence on psychology. It asserts a global materialism and also offers particular theories about how mind works. Its dual agenda is to build "machines that think" and to use machines to think about thinking. Its methodological premise is that if one builds a machine that can do something intelligent, the way one gets the machine to do it is relevant to thinking about how people do it as well.

Artificial intelligence is usually seen as having its strongest affinities with rationalist philosophies, defining knowledge as information and devaluing ambiguity when it offers a view of mind as program. But this view, commonplace within the literary culture, is only partial. One might call it the literary stereotype of the field. But AI has other dimensions that give it a far wider range of intellectual connections and implications for psychology.

This essay is interpretive, attempting to identify actual trends and influences of the computer on psychological thinking, and it is speculative, predicting a new alliance between AI and psychoanalysis, the latter being an intellectual companion far removed from rationalism, quantification, and formal propositions. This does not mean that there is an identity of spirit between computational and psychoanalytic models or even that there are not fundamental incompatibilities between them. But in addition to a list of traditional affinities between psychoanalysis and AI, they have something new in common. In recent years, computer scientists and psychoanalysts have talked a strikingly similar language about inner agents that construct the thinking and feeling self. Behind the language are shared concerns that suggest new theoretical linkages and a new source of vitality for psychoanalytic ideas.

Predictions are always dangerous. The point of making one here is what is gained from thinking it through: insight into how the computer's presence can act as a sustaining myth to support not one psychological culture but a range of them.

TRADITIONAL AFFINITIES

The very idea of AI—to create mind in machines—subverts traditional notions of the autonomous self in a way that parallels the

psychoanalytic enterprise. Most people see the autonomous self as an unproblematic idea because they have a day-to-day experience of having one. Our everyday language captures that experience and expresses the idea of free will; we say, "I act," "I do," "I desire." And even when people have learned through theology or philosophy to question the idea of free will, what they tend to do is make small modifications in their notion of the autonomous self; it becomes a self whose decisions are constrained. Inherent in psychoanalysis is a more radical doubt. The unconscious does not constrain; it constitutes a decentered self. Inherent in AI is an even more threatening challenge: If mind is program, where is the self? It puts into question not only whether the self is free, but whether there is one at all.

Traditional humanism is committed to the notion of an acting, intentional subject. In its challenge to the humanistic subject, AI is subversive in a way that takes it out of the company of rationalism and puts it into the company of psychoanalysis and radical philosophical schools such as deconstructionism. The psychoanalytic subject is decentered in the web of the unconscious; the deconstructionist subject is decentered in language; the computational subject is decentered—indeed, perhaps dissolved—in the idea of program.

These affinities will not reassure the traditional humanist who has gotten used to seeing AI as an enemy. They do not make AI any less of an assault on the idea of the self. But the attack comes from the left, so to speak, rather than from the right. Artificial intelligence is to be feared as are Freud and Derrida, not as are Skinner and Carnap.

The computational "explanation" of the chess move points to another way in which AI is more like Freud than Skinner. Within traditional science, and certainly for behaviorism, the line between subject and object is taken as sacred. But for Freud, his self-analysis, his technique of self-understanding, was indissociable from the development of his general theory. Like psychoanalysts, AI theorists have made a profession of dissolving the line between subjective and objective reflection. The intelligence embodied in the chess move is intelligence derived from personal knowledge of chess. "There's only one place to get ideas about intelligence, and that's from thinking about myself," says AI scientist Roger Schank. "In the end, I have just myself, and if it feels right, that's what I have to trust," says Donald Norman, an AI-influenced cognitive psychologist.⁴

Marvin Minsky, one of the founders and theoretical leaders of AI for the past quarter of a century, has always made it clear that as far as he is concerned, you can make a machine do only what you yourself know how to do. In order to build a program, you have to engage in self-analytic activity. In the early 1960s Minsky worked with a student, Thomas Evans, on an AI program that could pass the familiar visual-analogy tests: *A* is to *B* as *C* is to *D*, *E*, or *F*, where each letter stands for a geometrical drawing. His method was psychological: Think about yourself. And its reference point was psychoanalysis: "What you had to do was something like what Freud did. Tom Evans and I asked ourselves, in depth, what we did to solve problems like this, and that seemed to work out pretty well."⁵

Behaviorism rigorously forbids any reference to personal experience, and most other psychological schools try to ignore the issue. But AI and psychoanalysis have each articulated the need to integrate personal reference into theoretical construction. Each, in its own way, is a science of self-reflection.

But are such affinities superficial? After all, psychoanalysis explores the mind to discover the irrational; artificial intelligence invents machines through the exploitation of the rational. In fact, what stands between psychoanalysis and AI is not AI's "materialism." In the past quarter of a century, psychoanalysts have learned the necessity and the productivity of an intensified dialogue with psychopharmacology and neuroscience. And Freud himself hoped that someday his science of mind would be tied to its physical substrate, even if his own first efforts to make the connection had led to an impasse.⁶ What stands between psychoanalysis and AI is the view that AI is synonymous with rationalism, or rather with the kind of rationalism embodied in the idea of information processing.

If AI has seemed somewhat unitary in its implications for thinking about people, it is because what many observers know as AI is really information processing, a rule-driven, hierarchical approach to creating intelligence. But information processing is only one part of a larger picture.

THE TWO AIs

In the mid-nineteenth century George Boole formalized rules of logical inference in an algebraic form systematic enough that he felt

entitled to call his work "The Laws of Thought."⁷ Of course, Boole's title reached beyond his achievement, which is far from an all-inclusive model of mind. For one thing, Boole's laws need an external agent to operate them.

Boole's laws are something a person could use, but a computational version of Boole breathes life into his equations. An operator in the form of a computer program is placed within the system. Once there, the operator and the laws can be seen as a functioning model, if not of the mind, at least of a part of the mind.

One major branch of AI research can be described as doing just this—pursuing Boole's project in computational form. Information-processing AI gives active shape to formal propositions to create an embodiment of intelligence as rules and reason. Boole formulated algebraic rules for the transformation of logical propositions. Modern computer science has enlarged the logical and propositional to a more general notion of what it calls information, and it has enlarged algebraic transformation to a more general notion of computational processing. Boole would recognize a kinship between his project and Newell and Simon's way of putting these two advances together in GPS and other programs that laid the foundation for information-processing AI.

But artificial intelligence is not a unitary enterprise. Computation is a stuff out of which many theories can be fashioned. It is true to say that there is not one AI but many. And it is helpful to say that there are essentially two. The first is information processing, its roots in logic, the manipulation of propositions to obtain new propositions, the combination of concepts to obtain new concepts. The second comes from a very different style of work, present from the earliest days of the field but now having increasing influence, to the point of being the focus of attention wherever AI is discussed, from research seminars to popular articles. This second is "emergent AI."

Emergent AI has not been inspired by the orderly terrain of logic. The ideas about machine intelligence that it puts forward are not so much about teaching the computer as about allowing the machine to learn. This AI does not suggest that the computer be given rules to follow but tries to set up a system of independent elements within a computer from whose interactions intelligence is expected to emerge. From this perspective, a rule is not something you give to a computer but a pattern you infer when you observe the machine's behavior,

much as you would observe a person's. Its sustaining images are not drawn from the logical but from the biological.

Information processing breathes life into Boole by putting an operator into his system, but what it operates on shares the static nature of Boole's propositions. In traditional computers, millions of units of information sit in memory doing nothing as they wait for the central processor to act on them, one at a time. Impatient with this limitation, the goal of emergent AI is "pure" computation. Here, the whole system is dynamic, with no distinction between processors and the information they process. Families of neuronlike entities, societies of anthropomorphic subminds, and sub-subminds are in simultaneous interaction. The goal, no less mythic than the creation of a strand of DNA, is the generation of a fragment of mind.

The two AIs, rule-driven and emergent, logical and biological in their aesthetic, fuel very different fantasies of how to build mind from machine. If information-processing AI is captured by the image of the knowledge engineer, hungry for rules, debriefing a human expert in order to embody that expert's methods in algorithms and hardware, emergent AI is captured in the image of the computer scientist, up all night watching the twinkling lights of a computer in the hope that the interaction of "agents" within the machine will create intelligence.

Widely associated with the spirit and substance of the field as a whole (here I have called it the literary stereotype), information processing put AI in a distant relationship to psychoanalysis, whose ideas do not easily translate into rules or algorithms.⁸ Indeed, I now turn to how popular notions about AI drawn from information processing suggest that AI is all the things that psychoanalysis is not. My thesis follows directly: when the stuff of AI is expanded to include not only information but also active and interactive inner agents, there is a starting place for a new dialogue between the psychoanalytic and the computer cultures.

INFORMATION PROCESSING AND PSYCHOANALYSIS

The Freudian slip is a tempting target for psychologists bent on finding computerlike mechanisms behind human behavior. After all, one understands only too well the kinds of errors computers make. What sort of computer would make the kind of error that Sigmund

Freud saw as revealing a very different kind of meaning? In other words, what kind of computers are we?

In *The Psychopathology of Everyday Life*, Freud discusses slips of the tongue and takes as one of his examples a chairman who opens a parliamentary session by declaring it closed. The Freudian interpretation of this slip focuses on the complex feelings that may lie behind. Is the chairman anxious about the session? Does he have reason to believe that nothing good will come of it? Would he rather be at home? The slip is presumed to tell us about real wishes. Its analysis lays bare the concept of ambivalence—in this case, the chairman's mixed emotions about attending the session at all.

How can we see this human slip as an information-processing error? An MIT computer science student had no trouble finding an explanation: "A bit was dropped—the sign bit. There might have been a power surge. No problem." It's interesting that Freud saw a problem precisely because *open* and *closed* are so far apart—their opposite meanings give significance to their substitution. For computer science students, *open* and *closed* are close together. In their conceptual world, it is natural to code opposite concepts as the same root with a different "sign bit" attached (hot = -cold, dry = -wet, open = -closed). So if you think of the human mind as storing information in a computer's memory, substituting *closed* for *open* is easily justified. It might have been a small technical failure due to something as trivial as a power surge. It needs no recourse to the idea of ambivalence, hidden wishes, or emotional conflicts. What was interpreted in terms of sexually charged feelings, as a window onto conflicts, history, and significant relationships, becomes a bit of information lost or a program "derailed." What psychoanalysis would interpret in terms of *meaning*, this computational psychology would see in terms of *mechanism*.⁹

There is another way to look at the difference between the psychoanalytic and the information-processing view of the slip, a perspective that looks at the "width of determinism" in a system of interpretation. As a way of knowing, psychoanalysis has a logic that calls the whole person into play to explain all of his or her actions. This is why an individual can use something as small as a verbal slip to get in touch with the deepest levels of personality. What places the student's saying "There might have been a power surge. No problem" in such radical conflict with psychoanalysis is not so much that

power surging is alien to psychoanalytic categories but the idea that any single factor could explain an act of language.

In traditional logic, when you say, "All men are mortal; Socrates is a man; therefore Socrates is mortal," your conclusion is determined by two premises. Change one, and you get a new conclusion. Similarly, with an information-processing computer model, you drop one bit, one piece of information, and you get a new output. The determination is "narrow," like a highway with one lane. Psychoanalysis uses "wide" determination. It is based on another kind of logic, more like the logic that leads you to say that Shakespeare is a great poet. Coming across a bad poem by Shakespeare does not call the statement into question. Nor would the discovery that several of Shakespeare's best poems were written by someone else. So even if the chairman announced that the meeting was closed in the context of his wife being ill, her illness and his desire to be at home would not determine his slip in any simple sense. Psychoanalytic phenomena are as "overdetermined" as judgments of literary merit. Although popular images of a psychoanalytic dream book abound—along with a history of popularizers attempting to write one—there is no such thing as a "look-it-up" dictionary of Freudian symbols. The meaning of a dream can only be deciphered from the complex fabric of a particular dreamer's associations to it.

But computation is not synonymous with the narrow determination of information processing. Emergent AI builds models with broader determination. Whereas information processing gives concepts like *closed* and *open* actual symbolic representation in a computer, the building blocks of emergent AI do not have that kind of one-to-one relationship with such ideas. In symbolic representation, knowledge is stored as a static copy of a pattern. In an emergent system, the pattern itself is not stored. What is stored is data about the relationships among agents that would be expected to recreate the pattern. In this kind of system it is not possible for "one bit dropped" or "one rule changed" to make a difference to an outcome. In emergent systems, probabilities take the place of algorithms; statistics take the place of rules.

In a memoir she wrote in 1842, Lady Ada Lovelace, a friend and patroness of Charles Babbage, inventor of the "analytical engine," was the first person to go on record with a variant of the oft-quoted statement "Computers only do what you tell them to do."¹⁰ The

Lovelace model for thinking about computers' strengths and limitations is paradigmatic for information processing. But it does not hold for emergent AI. Here, the point is quite precisely to make computers do more than they were ever told how to. It has become commonplace for people to quote Lovelace to defend against the idea that we are like machines: "People are not computers. They don't follow rules. They learn. They grow." But emergent AI is characterized by "anti-Lovelace" representations of the computer. It breaks down resistance to seeing continuity between computers and people by describing computers that learn and grow, by describing computers whose resonance is biological rather than logical.

EMERGENT AI AND BROAD DETERMINATION

This biological resonance is illustrated by the perceptron, a pattern-recognition machine designed in the late 1950s and a good first example of emergent AI. Information-processing AI is made out of data and rules. Emergent AI is made out of very different stuff, a stuff most easily captured in anthropomorphic language.

Imagine that you have access to the opinions of a thousand simple-minded meteorologists, each of whom has a different unreliable method of weather forecasting. Each bases a judgment on a fragment of evidence that may or may not be related to predicting rain. How do you form a judgment? A narrowly determined method, in an information-processing system, for example, might be autocratic—identifying the meteorologist with the best track record and going with that vote. Another strategy, both more democratic and more broadly determined, would be to let the majority decide. The perceptron refines the democratic strategy by weighing each vote with a number related to the individual meteorologist's past record.

So, for example, to get a perceptron to recognize a triangle, you show it samples of triangles and nontriangles and make the system "guess." Its first guesses are random. But the perceptron is able to take advantage of signals saying whether its guess is right or wrong to create a voting system in which agents who have guessed right get more weight. Perceptrons are not programmed, but learn from the consequences of their actions.

In the narrowly determined method, you would have complete breakdown if the chosen meteorologist went insane. But in the brain,

damage seldom leads to complete breakdown. More often it produces a degradation of performance proportional to its extent. In other words, when things go wrong, the system still works, but not as well as before. Information-processing systems lose credibility as models of mind because they lack this feature; the perceptron shows the graceful degradation of performance typical of the brain. Even with some disabled meteorologists on board, the perceptron still produces the best possible decision based on the subset of functioning actors.

In an information-processing model, intelligent behavior follows from fixed rules. In the perceptron there are none. There is no flow-chart, no rule-driven path through the system. Nor are there one-to-one correspondences between information and output. What is important is not what an agent knows but its place in a network, its interactions and connections. The perceptron presents a model of mind as a society in which intelligence grows from the cacophony of competing voices.

In an information-processing model, the concept "rain" would be explicitly represented in the system. In the perceptron the decision "it will rain" is born from interactions among agents, none of whom has a formal concept of rain. Perceptrons show the emergence of what information processing takes as its raw material. Information processing begins with formal symbols. Perceptrons, like Freud's unconscious, operate on a subsymbolic and a subformal level. And most important for the current discussion, perceptrons rely on the interactions of inner agents, objects within the system.

Object theory is a central aspect of emergent AI and forms the link between AI and new directions of psychoanalytic thought. The inner agents in perceptrons make a bridge to the broad determinism of psychoanalysis. But it is only an opening. After perceptrons and the perceptronlike systems of the 1960s, it took another round in the development of computational ideas before inner objects came to occupy center stage. This is the story to which I now turn—the story of a second generation of emergent AI with an emphasis on inner objects and a new pathway for influence on psychoanalysis.

EMERGENT AI AND COMPUTATIONAL OBJECTS

The atmosphere in the AI laboratories of the early 1960s was heady. The work of Norbert Wiener, John von Neumann, and Alan Turing

had set off shock waves that were still fresh. The first information-processing programs that emulated fragments of human thought had only recently produced their surprise. Perceptronlike models (and there were many of them, including Oliver Selfridge's "Pandemonium" and Warren McCulloch's "neural nets") led researchers to biologically resonant descriptions of artificial mind. Thoughts were on the ultimate nature of intelligence.

Artificial intelligence researchers saw little reason for a more humble style. On the contrary, AI defined itself as an enterprise of mythic proportions: mind creating mind. In doing so, the field drew a certain kind of person into its culture, not unlike the kind of person drawn into the early circle around Freud. There, too, the enterprise was mythic: the rational understanding the irrational. There, too, it was without precedent or academic security. The first generation of AI researchers, with backgrounds as diverse as mathematics, psychology, economics, and physics, like the first generation of psychoanalysts, had not been trained in "the field" because it did not exist. There was no academic discipline. There were only new worlds to conquer.

In the early 1960s emergent models were as much a part of what seemed exciting in AI as information-processing programs. But for almost a quarter of a century, emergent AI seemed swept aside. In its influence on psychology, AI became almost synonymous with information processing. Newell and Simon developed rule-based systems in purest form—systems that simulated the behavior of people working on a variety of logical problems. Such simulations offered the promise of more—the promise of making artificial mind out of rules. And if you can build mind from rules, then mind can be presumed to have had rules all along. Following this logic, researchers made information-processing models the backbone of cognitive science.

The language of information processing—descriptions of "search," "subroutine," "scripts," and "grammars"—became common currency among psychologists who accepted the idea that "toy programs," little pieces of machine-embodied intelligence, were representative of bigger things to come. Computer programs that could play chess, manipulate blocks, or "converse" with imaginary waiters in imaginary restaurants did more than model small pieces of mental functioning. They supported the idea that the means used to build

them, all drawn from the information-processing paradigm, might someday capture the essence of mind. This idea was bolstered by the wordly success of a particular kind of information-processing program—the expert system. In it the AI scientist extracts decision rules from a virtuoso in a field (medical diagnosis, for example) and embeds them in a machine that will then do the diagnosis “for itself.”

By the mid-1970s AI was no longer marginal. It had its own academic programs, its own journals, its own conferences. It was well funded because of its value in the marketplace and to the military. Expert systems were used to analyze stock prices, data from oil well drillings, materials from chemical samples. Companies competed to hire AI graduates to start in-house departments. The future of the field became part of a heated discussion about Japanese-American industrial rivalry.

Now AI could promise a more traditional kind of career, much as the medicalization of psychoanalysis paved the way for it to become a professionalized psychiatric specialty. In both psychoanalysis and AI, traditional careers meant new pressure to engage in the kind of work that promised visible results. In psychoanalysis the pressure was to “cure,” to work on educational problems, to do “applied psychoanalysis.” In AI research the pendulum swung from what had been most mythic about the dreams of the 1950s and early 1960s to what people “knew how to do”—gather rules and code them in computer programs.

But even as the information-processing model reached near-hegemony in the late 1970s, the conditions for something very different were developing. First, there was important technical progress. Computer scientists had long strained against the limitations of the von Neumann computer, in which one processor might manipulate the passive data in a million cells of memory. It had always been obvious that, in principle, the distinction between processor and memory could be abolished by making every cell in the computer an active processor. Doing so, however, had always been prohibitively expensive. But now, projects such as the Connection Machine were realistic enough to be funded. There, the plan is to have a million microprocessors put together to make one computer whose memory and computational power are fully distributed. No longer would there be an operator and the passive material it operated on. Computation was “waking up from the Boolean dream.”¹¹

Along with hardware that presented fresh possibilities were new ideas about how to program it. The development of programming methodologies with suggestive names like "message passing" and "actor models" created the context for thinking about computational agents in communication. Standard computer programs are lists of instructions in the form of imperatives: "add these numbers," "put the result in memory," "get the content of that memory location." Artificial intelligence programs in LISP or Prolog operate on more abstract data but still consist of instructions for manipulating information. The first quarter of a century of the development of programming was based on a process language for describing how to pass information from one place to another. But researchers now felt the need to deal with a different kind of event: not the *passing* of something but the *making* of something. By a coincidence that turns out to be highly suggestive for the present discussion, computer scientists called their so-far most prominent response "object-oriented programming."

If you want to simulate a line of customers at a post office counter (in order to know, for example, how much longer the average wait would be if the number of clerks were to be reduced by one), you write a program that creates an internal object that "behaves like" a person in a line at the post office. It advances when the person ahead in the line advances; it knows when it has reached the counter and then proceeds to carry out its transaction. The contrast between this object-oriented approach and traditional programming strategies is dramatic. A traditional FORTRAN programmer would assign x 's and y 's to properties of the customers and write computer code to manipulate the variables. Object-oriented programming refers directly to the inner objects that represent the customers in line: x 's and y 's do not appear.

In object-oriented programming, the programmer makes new objects that, once created, can be "set free" to interact according to the natures with which they have been endowed. The programmer does not specify what the objects will actually do, but rather "who they are."

If something of the "feel" of an information-processing program is captured by the image of the flow chart, something of the "feel" of object-oriented programming is captured by the pictures of file folders, scissors, and wastebasket that appear on the screens of

computers with an iconic interface. The icons are a surface reflection of a programming philosophy in which computers are thought of as "electronic puppet shows" and "there are no important limitations to the kind of plays that can be enacted on their screens, nor to the range of costumes or roles that the actors can assume."¹² For mathematicians, the algebraic manipulations in traditional programming have a compelling reality. But for most nonmathematicians, the object-oriented approach has a more direct appeal, the appeal of actors on a stage.

By the early 1980s the coexistence of new parallel hardware and new ideas about objects in programming set the stage for the pendulum to swing away from information processing. The beginning of the decade saw the first of a growing series of papers from very different origins—engineers eager to build new parallel machines, computer scientists eager to try mathematical ideas that could guide new efforts at parallel programming, psychologists looking for new models that had a biological (indeed, a neurological) resonance. Emergent AI had not so much died as gone underground. It reemerged with a vengeance and with a new label: "connectionism." Once again, there would be no distinction between the processor and what it processed. There would be no specified set of operations. There would only be communities of agents in direct interaction with each other.

But proponents of the new theory of connectionism go beyond earlier stages of emergent AI in the steps they want to take away from Boole. For example, the perceptron could not itself generate new objects or elucidate how new objects could emerge. Its agents were programmed by a human acting from the outside. Today's connectionists hope to go further by bringing together parallel machines, the maturation of ideas about how to program them, and most important, a new sense of the central problem facing the field, something that had scarcely been formulated during the 1960s: How are objects created?

PSYCHOANALYTIC OBJECTS AND SOCIETY MODELS

In the focus on inner objects and their emergence and interaction, AI shares preoccupations that are central to contemporary psychoanalytic theory. As was the case in AI, the development of a psychoanalytic object theory is a later development of the field. It was not where the theory began.

Early psychoanalytic theory was built around the concept of drive, demand that is generated by the body and that provides the energy and goals for all mental activity. But later, when Freud turned his attention to the ego's relations to the external world, the significance and structure of these relations could not easily be framed in drive theory.

By 1917 Freud began to formulate a language to handle these matters. He described a process by which people form inner "objects." In *Mourning and Melancholia*, Freud argued that the sufferings of a melancholic arise from mutual reproaches between the self and an internalized father with whom the self identifies. In this paper Freud described the "taking in" of people (in psychoanalytic parlance, objects, and in this case the father) as part of a pathology, but he later came to the conclusion that this process is part of normal development. Indeed, this is the mechanism for the development of the superego—the taking in, or introjection, of the ideal parent.

According to Freud, we internalize objects because our instincts impel us to. In his work the concept of inner objects needed to coexist with the scaffolding of drive theory. But many psychoanalytic theorists who followed him were less wedded to the drive model. They widened the scope of what Freud meant by "object relations" to the point where we now think of them as a distinctive school. Classical Freudian theory has many overlapping concepts to describe internal objects: memory traces, mental representations, introjects, identifications, and the idea of inner structures such as the superego. The object relations approach is more specific about what we contain. It describes a society of inner agents, or "microminds"—"unconscious suborganizations of the ego capable of generating meaning and experience, i.e. capable of thought, feeling and perception."¹³ Relationships with people, "brought inside" as inner entities, are the fundamental building blocks of mental life.

Whereas Freud focused his attention on a single internalized object, the superego, object relations theorists described a richly populated inner world. Psychoanalyst Melanie Klein went so far as to characterize the people that the child brings inside (as well as the representations of parts of the body) as having psychological features, personalities. They can be seen as loving, hating, greedy, envious. Psychoanalyst W. R. D. Fairbairn even reframed the basic Freudian motor for personality development in object relations terms. For

Fairbairn, the human organism is not moved forward by Freud's pleasure principle, the desire to reduce drive tension, but rather by its need to form relationships. This constitutes a profound recasting of the psychoanalytic view of the self: people are not fundamentally pleasure-seeking, but object-seeking.

The language that psychoanalysts need to talk about objects—how they are formed, how they interact—is very different from the language they need to talk about drives. In his "Project for a Scientific Psychology," Freud tried to use informationlike terms derived from the description of the reflex arc—the pain fiber carries information to the brain, the motor fiber carries information to the muscle—to talk about memory, instincts, and the flow of psychic energy. But information metaphors break down completely when you use them to talk about inner objects. As in object-oriented programming in computer science, so it is in psychoanalysis. When one talks about objects, the natural metaphors have to do with making something, not carrying something.

In classical psychoanalytic theory a few powerful inner structures—the superego, for example—act on memories, thoughts, and wishes. Object relations theory posits a dynamic system in which the distinction between processor and processed breaks down. The parallel with computation is clear: in both cases there is movement away from a situation in which a few inner structures act on a more passive stuff. Fairbairn replaced the Freudian dichotomies of ego and id, structure and energy with independent agencies within the mind that think, wish, and generate meaning in interaction with each other, much as emergent AI sets free autonomous agents within a computer system.

The development of object relations theory has led psychoanalysts to ask if allegiance to Freud depends on accepting his drive model. Some have tried to preserve Freud's original drive language but to use it in a way that accommodates a new emphasis on object relations—for example, by assigning objects a role in relation to the discharge of drive: they may inhibit, discharge, facilitate, or serve as drive's target. But this reworking of language is less a solution than an attempt to gloss over the problem. It only works if inner objects do not have elaborate properties or if their creation is seen as an occasional event.

But when objects become central to one's understanding of the psyche there is greater pressure to move away from drive theory.

Although drive theory has become increasingly sophisticated and open to the discussion of inner objects, the split between a drive approach and an object relations approach is a central division in psychoanalysis today.¹⁴ The division is parallel to the split between information processing and emergent AI. To use Thomas Kuhn's language, object relations theorists are saying that psychoanalysis can no longer proceed as "normal science,"¹⁵ growing by the assimilation of new data into the old theory. For them, object relations is a paradigm shift within psychoanalysis, much as the hypothesis of emergence—that intelligence grows out of the interaction of multiple agents (it is not what you know but who and where you are)—represents a paradigm shift in AI.

Artificial intelligence theorists Marvin Minsky and Seymour Papert have built a computational model that evokes Fairbairn's object relations theory. Their model takes the mind as a society of interacting agents. These agents are anthropomorphized, discussed in the terms one usually reserves for a whole person, but they do not have the complexity of people. Indeed, their model is based (as was the perceptron concept) on these agents being "dumb." Each knows one thing and one thing only. And, like the "voting agents" in the perceptron, their narrowness of vision leads them to very different opinions. The complex structure of behavior or emotion or thought emerges from the conflict of their opposing views.

The most elaborate presentation of this theory, Minsky's book *The Society of Mind*, describes a vast array of agents: censor agents, recognition agents, and anger agents, to name only a few.¹⁶ Not surprisingly, Minsky recognizes Freud, who also wrote extensively about censor agents, as a colleague in "society" modeling. More surprisingly, Minsky sees censors as key actors, not only for modeling human thought but also for making intelligent machines.

Minsky's idea of the censor is a dramatic example of the developing resonance between psychoanalysis and emergent AI. Freud's censor protects people from painful thoughts. The extension of this idea to cognitive functioning and to the "thoughts" in a machine does not depend on the assumption that the agents or the system as a whole feels pain. To function coherently, according to Minsky, an intelligent system must develop a certain inattention to its contradictory agent voices. Minsky's formulation is that there cannot be intelligence, artificial or otherwise, without repression. Allen Newell

has talked about the necessity for censors in large and complex information-processing systems. But with clear, unambiguous rules stated in advance, an information-processing computer can also do without them. Censors may turn out to be practical, but they are not theoretical necessities in the information-processing paradigm. In the case of society theory, however, censors are intrinsic. Since there cannot be intelligence without contradiction and conflict, only the presence of censors allows intelligence to emerge.

In this, society and Freudian theory join on an important point. Freud did not "discover" the unconscious. His contribution was the elaboration of a *dynamic* unconscious. What is unconscious is not simply forgotten, old, or irrelevant to current functioning. It is repressed. Powerful forces keep it down, and for good reason. Similarly, for Minsky, what is repressed in the computational machine and in what he has called the human "meat machine" *needs* to be repressed.

Freud wrote about the effects of the repression of frightening, emotion-laden experience. Minsky extends Freud's ideas to the cognitive domain. "A thinking child's mind . . . [needs no one] to tell it when some paradox engulfs and whirls it into a cyclone." Paradox, argues Minsky, is as dangerous as the primal scene. The child knows it is in the presence of a threat when it is asked to sketch the nonexistent boundaries between the oceans and the seas or to consider questions about the chicken and the egg, about what came before the start of time, and about where the edge of space is. Minsky adds: "And what of sentences like '*This statement is false,*' which can throw the mind into a spin? I don't know anyone who recalls such incidents as frightening. But then, as Freud might say, this very fact could be a hint that the area is subject to censorship."¹⁷

Minsky feels that the notions of "cognitive repression" and the "cognitive unconscious" will allow us to go beyond Freud. He uses Freud's discussion of jokes as an example. Freud's 1905 work on jokes explained that inner censors serve as barriers against forbidden thoughts. Most jokes are stories designed to fool the censors. It is a way to enjoy a prohibited wish. This is why so many jokes involve taboos concerning cruelty and sexuality. But it troubled Freud that this theory did not easily account for "nonsense jokes." One of Freud's hypotheses about the power of the nonsense joke was that senselessness reflects "a wish to return to carefree childhood, when

one was permitted to think without any compulsion to be logical." The idea of the cognitive unconscious supports this view: paradox and senselessness need to be repressed in the process of developing emergent intelligence, whether in machines or in people. Absurd results of reasoning are taboo, as threatening as sex. The censors work as hard to suppress them; they have no innocence.

SUBVERSION AND NORMALIZATION

Despite their differences, psychoanalysis and AI have always shared theoretical affinities—among these, as we have seen, the challenge to the idea of the autonomous, intentional actor, the need for self-reference in theory building, and the need for objects such as censors to deal with internal conflict. But the affinity became something stronger when the cluster of issues about objects came to occupy center stage for both. This new orientation has made the old common elements more common: agent theories in AI highlight theoretical concerns that echo psychoanalytic ones. These include conflict, internal inconsistency, and perhaps most dramatically, the subversion of the subject, the "decentered" self.

Although both psychoanalysis and AI have always challenged the actor "I," both have theoretical variants that underscore this challenge more than others. The Freudian unconscious undermines the idea of a unified subject, but many of Freud's followers moved toward restoring a sense of there being a mental executive by concentrating their attention on the ego, that part of Freud's divided self that was turned outward toward reality. These "ego psychologists" began to talk about it as an agent capable of integrating the psyche. Anna Freud wrote of its powerful artillery, the mechanisms of defense. Heinz Hartmann argued that the ego had an aspect that was not tied up in the individual's neurosis, a "conflict-free" zone. Hartmann wrote about this unhampered aspect of the ego as though it were free to act and choose, independent of constraints. It almost seemed the seat for a reborn notion of the will, a locus of moral responsibility. Intellectual historian Russell Jacoby, writing of ego psychology's reborn, autonomous "I," went so far as to describe it as the "forgetting of psychoanalysis."¹⁸

In its subversive form, which splits the ego and undermines the subject, psychoanalysis is hard to take. It flies in the face of common-sense understanding. It is a subversive science. Ego psychology normalizes it. It takes what is most subversive—the decentered self—and softens it. Ego psychology presents the version of the unconscious most acceptable to the conscious.

This pattern of a normalizing response is common to all subversive sciences of mind, including, of course, computational ones. We saw how the very idea of AI calls the self into question through the notion of program. But AI, too, has variants that soften its message about decentering. For example, if you reduce AI to the idea of expert systems, it is a small step to think of the expert system as a resource on which some not clearly specified central executive can call. When you begin with the idea that a computer might have such an executive and such resources, the idea that a human has them too follows directly. It makes an AI model of mind seem less threatening because what needs to be thought of as computational and rule-driven, alienated from intention, is not my “I,” but my “expert” in a limited domain—for example, that part of me that can play chess. The self becomes the executive who oversees the expert. So there are versions of both AI and psychoanalysis that defuse the subversive decentering principle by restricting its role to explaining parts of the mind and thus avoid the risk of dissolving the whole.

This strategy for neutralizing subversive theory is less viable in the case of agent and object theories, which are more aggressive in their denial of a unified self. Indeed, these theories define themselves through that denial. They put psychoanalysis and AI in a new and closer relationship with each other and with other intellectual movements that “deconstruct” the humanistic subject.

The strength and the weakness of object theories are the same in both psychoanalysis and AI: the strength is a conceptual framework that offers rich possibilities for models of interactive process; the weakness is that the framework may be *too* rich. The postulated objects may be too powerful: they explain the mind by postulating many minds within it. Object theory confronts both fields with the problem of infinite regress. There is something deeply unsatisfying in a theory that cannot go beyond assuming a homunculus within the human, for how then do we explain the homunculus within without postulating yet another one within it, and so on?

Psychoanalytic theorists struggle with this issue. Within the field much of the criticism of overpowerful inner objects has taken Melanie Klein's work as its target. For example, psychoanalyst Roy Schafer has argued that Klein and the English School of object relations have carried the reification implicit in Freudian metapsychology to a "grotesque extreme": "A multitude of minds is introduced into a single psychic apparatus. . . . The person is being envisaged as a container of innumerable, independent microorganizations that are also microdynamisms."¹⁹ Essentially, Klein's critics feel that her idea of "inner idealized figures protecting the ego against terrifying ones is tantamount to proposing that there are internal friendly and hostile 'demons' operating within the mind."²⁰

Kleinians reply that these internal figures are not demons but unconscious fantasies and thoughts. They are the ideas we have about what we contain.²¹ But this response hardly settles the question. Psychoanalyst Thomas Ogden puts the problem starkly: How can thoughts behave as agents?

If internal objects are thoughts . . . then they cannot themselves think, perceive, or feel, nor can they protect or attack the ego. Even to the present, Kleinian theorists have not been able to disentangle themselves from the Scylla of demonology and the Charybdis of mixing incompatible levels of abstraction (i.e. active agencies and thoughts).²²

In computer science, connectionism has not solved the problem of accounting for objects (what they are and how they come into being). Connectionism simply postulates the inner agents it requires, which is why AI scientist Terry Winograd has gone as far as to say that part of its appeal is that "it has a higher percentage of wishful thinking."²³ But the problem of infinite regress (accounting for the entities that are then to account for thought) has a very different cast in AI than in psychoanalysis because computer scientists are used to relying on a controlled form of circular reasoning—"recursion"—as a powerful technical tool.

Most of us learned at school to define x to the power n as x multiplied by itself n times. Power is defined in terms of multiplication. Computer scientists prefer to define x to the power n as x to the power $n-1$ multiplied by x . Power is defined in terms of power. From such simple examples, which are shared with precomputational

mathematics, computer science has built a mathematical culture that relies heavily on defining things in terms of themselves.²⁴

For the psychoanalyst Ogden, the idea that a thought might think was unthinkable. Information-processing AI also divides thought from thinking. What is closest to a thought is information. What is closest to thinking is its processing. But emergent AI breaks down this distinction. It takes the idea of recursion and turns it into an overarching aesthetic. To put it more sharply, emergent AI provides a way out of the problem of infinite regress by redefining the problem as a source of power. Taking recursion as a scientific aesthetic gives AI a way out of a theoretical hole. It may offer a similar possibility to psychoanalysis.

What computational memory was to the birth of cognitive science in the 1950s, recursion could be to psychoanalytic studies of the 1990s. One could imagine computer scientists trying to support Kleinian psychoanalysts by building a detailed computer model of Kleinian objects. But one could also imagine computationally sophisticated psychoanalytic theorists finding, in the recursive idea that thoughts might think, a pleasing virtue rather than a devastating vice. One could imagine a psychoanalytic theorist seeing recursive ideas in his or her own work as a source of legitimation rather than a sign of weakness.

There can be no simple prediction about how recursion will help psychoanalysts deal with the infinite regress of object theories, but it seems probable that the kind of influence to look for is psychoanalysis becoming increasingly permeable to recursion as a sustaining myth. This would make the very thing that has been leveled as criticism become a way to support the theory. In the spirit of George Miller's account of computer memory and behaviorism, psychoanalysts might find it embarrassing to deny human thoughts the ability to think, when "computer thoughts" are presumed to do so.

If the central issue in psychoanalytic theory today turned on the nature of the "death instinct," there would be little that is helpful in theories about a machine that was never born. But to the degree that theoretical concerns in psychoanalysis have to do with the structure and functioning of internal objects, it is moving toward AI—to the point where the path to a productive dialogue seems open.

When the dialogue begins, the influence of AI on psychoanalysis will not necessarily be dependent on whether AI offers technical

advice but on whether it can offer moral support to beleaguered psychoanalytic object theorists in their debates. Can it serve as sustaining myth? The influence of AI on psychology, psychoanalytic and other, is related not only to the solution of technical problems but also to the growth of psychological cultures.

PSYCHOANALYTIC CULTURE AND COMPUTER CULTURE

Psychological cultures do not exist only in the world of professionals. Artificial intelligence and psychoanalysis set the context in which professional psychologists and the amateur psychologists we all are think about thinking. From a sociological perspective on this wider psychological culture, object theories make ideas in AI and psychoanalysis more "appropriable," easier for people to take up as ways of thinking about themselves, than theories about information or drive. In other words, object theories give psychoanalysis and AI a greater presence as philosophies in everyday life. Fairbairn's dense texts and the mathematical theory of connectionism might not be any more accessible to lay thinking than technical papers on information processing or on psychoanalytic drives. But when object-oriented theories are popularized and move out into the general culture, they have a special appeal. Ideas about objects and agents are more concrete than ideas about drives and flow-charts. They are seductive because it is easy to "play" with them. And they speak to a common problem. We all have the experience of not feeling completely "at one" with ourselves: inner voices offer conflicting advice, reassurance, and chastisement. These experiences are easily and satisfyingly translated into a drama of inner objects.

Freudian ideas about slips of the tongue became well known and gained wide acceptance for reasons that had little to do with positive assessments of their scientific validity. Freudian slips became part of the wider psychological culture because they made it easy to play with what might be hidden behind them. The slips are almost tangible ideas. They are manipulable. Slips are appealing as objects to think with. You can analyze your own slips and those of your friends. The theory of slips provided a way for psychoanalytic ideas to become part of everyday life. They helped to make psychoanalytic theory appropriable.

A Freudian perspective on the appropriability of psychoanalytic ideas might go further to suggest that the theory of significant slips is appealing because it puts us in immediate contact with the taboo. We are afraid of the sexual and aggressive sides of our natures, but we want to be in touch with them as well. Psychoanalytic ideas give us a way to play with what is forbidden. Similarly, we are afraid to think of ourselves as machines, yet we want to find a way to acknowledge this very real, if disturbing, part of our experience. Playing with AI, with the idea of the mind as computer, makes this possible. Now, playing with computational and psychoanalytic theories of objects and agents allows us to go even further. The idea of agents gives us a way to acknowledge the experience of fragmentation. The rational bias in our culture presents consistency and coherency as natural, but feelings of fragmentation abound. Indeed, it has been argued that they are a contemporary cultural malaise.²⁵ Theories within psychoanalysis and AI that speak simply and dramatically to the experience of a divided self have particular power.

In the past the computer culture and the psychoanalytic culture have been separate. In the main, psychoanalytic ideas for thinking about the self were congenial to people who had little contact with computational ones. If and when members of the psychoanalytic culture met computational models of mind, they were most likely to be information-processing models that seemed out of step with a psychoanalytic outlook. These models described sequences, not associations, and their model of determination was narrow rather than wide. But increasingly, the computational ideas put forward and reported in the popular, as well as the academic, press are not about rules and information but about agents, connections, and societies of mind. These new metaphors have a biological aesthetic—they are the kind of things that could be going on in a brain. They suggest broad determination and dynamic repression. They describe a system in conflict. And, most important, they resonate with the psychoanalytic ideas that are currently abroad, ideas not about drives and their vicissitudes but about objects and their interactions.

When the computer presence relegitimated the idea of memory, it was reinforcing an idea about psychology that predated computation. But ideas about recursion and agents are not precomputational. Dare one speculate what will pass between computation and our psychological culture if AI finds a voice finally divorced from what

was static in logic and if psychoanalysis finds a voice finally divorced from the issues of nineteenth-century drive theory?

ENDNOTES

- ¹Many of the ideas in this paper emerged in a series of conversations with Seymour Papert, a collaborator in the development of my notion of the role of sustaining myths in the sociology of the sciences of mind.
- ²Cited in Jonathan Miller, *States of Mind* (New York: Pantheon, 1983), 23.
- ³Ibid.
- ⁴Cited in Sherry Turkle, *The Second Self: Computers and the Human Spirit* (New York: Simon and Schuster, 1984), 256.
- ⁵Cited in Jeremy Bernstein, *Science Observed* (New York: Basic Books, 1982), 110–11.
- ⁶Sigmund Freud, "Project for a Scientific Psychology," *The Standard Edition of the Complete Psychological Works of Sigmund Freud*, vol. 1, trans. and ed. James Strachey (London: Hogarth Press, 1960).
- ⁷George Boole, *The Laws of Thought*, vol. 2 of *His Collected Works* (La Salle, Ill.: Open Court Publishing Company, 1952).
- ⁸A suggestive effort to construct psychoanalytic algorithms was made by French psychoanalyst Jacques Lacan in his theory of the *mathèmes*. The power of this idea derives from its effort to legitimate systematicity and a closer relationship with science in psychoanalytic studies. See Sherry Turkle, *Psychoanalytic Politics: Freud's French Revolution* (New York: Basic Books, 1978.)
- ⁹For an example of an information-processing perspective on the Freudian, see Donald Norman, "Post-Freudian Slips," *Psychology Today*, April 1980:41–44ff; Norman, *Slips of the Mind and an Outline of a Theory of Action* (San Diego: Center for Human Information Processing, University of California, November 1979); and Norman, "Categorization of Action Slips," *Psychological Review* 88 (January 1981):1–15.
- ¹⁰Lovelace put it like this: "The analytical Engine has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform."
- ¹¹This phrase is borrowed from Douglas R. Hofstadter, who discusses computation and the Boolean aesthetic in "Waking Up From the Boolean Dream, or Subcognition as Computation," in *Metamagical Themas: Questing for the Essence of Mind and Pattern* (New York: Basic Books, 1985).
- ¹²Alan Kay, "Software's Second Act," *Science* 85 (November 1985):122.
- ¹³Thomas H. Ogden, "The Concept of Internal Object Relations," *The International Journal of Psycho-Analysis* 64 (1983):227.
- ¹⁴See Jay R. Greenberg and Stephen A. Mitchell, *Object Relations in Psychoanalytic Theory* (Cambridge: Harvard University Press, 1983).
- ¹⁵Thomas Kuhn, *The Structure of Scientific Revolutions*, 2d ed. (Chicago: University of Chicago Press, 1970).
- ¹⁶Marvin Minsky, *The Society of Mind* (New York: Simon and Schuster, 1987).
- ¹⁷Ibid., 183. Fieldwork with children and computers is rich in examples of the kind of fright that Minsky expects. For example, an incident where it was evoked by a first contact with recursion is reported in Sherry Turkle, *The Second Self*. Interviews with adults on early experiences also reveal many such memories—fear

of prisms, of mirrors reflecting mirrors, fear of questions such as "How far away are the stars?"

¹⁸Russell Jacoby, *Social Amnesia: A Critique of Contemporary Psychology from Adler to Laing* (Boston: Beacon Press, 1975).

¹⁹Roy Schafer, *A New Language for Psychoanalysis* (New Haven: Yale University Press, 1976), 3; and Schafer, *Aspects of Internalization* (New York: International University Press, 1968), 62.

²⁰Ogden, "Internal Object Relations," 229.

²¹Hannah Segal, *Introduction to the Work of Melanie Klein* (London: Hogarth Press, 1978).

²²Ogden, "Internal Object Relations," 230.

²³*Science* 86 (May 1986):27.

²⁴The computational aesthetic of recursive thought has been expressed in a poetic and accessible form by Douglas R. Hofstadter, who presents recursive phenomena as a source of power in Bach's music and Escher's art as well as in Gödel's mathematics. See *Gödel, Escher, Bach: An Eternal Golden Braid* (New York: Basic Books, 1978).

²⁵See, for example, Christopher Lasch, *The Culture of Narcissism* (New York: Norton, 1979).