

Reviews of Books

States of Mind. JONATHAN MILLER. Pantheon Books, 1983. 316 pp

Jonathan Miller has just completed a book, called *States of Mind*, that is based on a new television series shown on BBC, but not yet in America. The book is a very well edited transcription of fifteen interviews with psychologists, anthropologists, and sociologists, including such notables as George Miller, Jerome Bruner, and Rom Harré. The contributors probably familiar to most AI researchers are Daniel Dennett and Jerome Fodor, as well as two contributors well-known for their writing on art and perception, Ernst Gombrich and Richard Gregory.

The interviews are uniformly intelligent, original, and stimulating. As summaries of basic arguments about mental models, perception, and ethical questions of mental problems, you can't do better than this collection. Miller is especially good at explaining the motivations for traditional psychology and how it has evolved into its present-day "cognitive" form. The anthropologists and psychiatrists in particular are unrelenting humanists who provide convincing tutorials on how their fields have advanced, making previous work understandable without being condescending.

Probably the most interesting thing to be discovered in this book is how the simple theme of "constructive understanding" plays a pivotal part in modern explanations of perception, language, problem solving, and interpersonal relations. By this perspective the specialists share a common methodology for analyzing processes and systems that we are all intricately part of. The clarifications by Harré and Geertz, for example, of how language places us apart from animals, are stunningly clear and valuable—somehow obvious, but again and again missed by the public mind.

A collection of this kind provides a valuable opportunity for AI researchers to develop their models of reasoning by constructing interfield analogies. The studies in *States of Mind* fall naturally into two categories: mental/perceptive and social/psychiatric—micro and macro perspectives on the nature of thought. While it is common to suppose that AI needn't be limited to mechanisms used by the brain, we more rarely consider the large-scale implications of placing an independent intelligent agent in an uncertain world, an agent with perhaps contrary goals, whose expectations are sometimes violated, and whose failures to understand the world lead to continuous striving for an overarching meaning. For mankind, this is the stuff of psychology, religion, and magic. Machines may be freed of the limitations of the human brain, but how will they cope with the existential problems of the universe?

The overall scope of the book is quite broad; applying the ideas to AI requires some condensation and reinterpretation. Miller isn't writing for AI researchers, nor does he

fully realize in his commentary the striking similarities of methods and results in the diverse fields. Therefore, to synthesize what the contributors are saying for an AI audience, I'm going to develop the main theme of "constructive understanding" in some detail. This idea is so powerful and important—one of the key insights of our century—every AI researcher should know it as well as theories of logic or search. To whet your appetite, the discussion will show how Harré's analysis of violence in football games is related to Gombrich's study of drawing (and why AI researchers should care). We begin with a brief consideration of the evolution of cognitive psychology and the study of the brain.

Miller tells us that, "In its understandable effort to be regarded as one of the natural sciences, psychology paid the unnecessarily high price of setting aside any consideration of consciousness and purpose, in the belief that such concepts would plunge the subject back into a swamp of metaphysical idealism." (p. 32) In their interviews, George Miller and Bruner relate how theories of machines, such as servomechanisms and signal detection and information theory, reintroduced the idea of internal "states," something represented in the machine about its world. In this way, the use of terms like "goals" and "expectations" became respectable again—"the kinds of things (engineers) say about machines, a psychologist should be able to say about a human being." (p. 23) I have simplified the arguments here for brevity. The interviews provide a readable, technical discussion that might be expected in any advanced text.

The discussions by Gregory, Fodor, and Geschwind are full of fascinating examples. We are treated to a discussion of our inability to intellectually juggle an optical illusion, raising intriguing questions about the effects of biased expectation on higher problem solving. Fodor's discussion of a Wittgenstein thought experiment fairly demolishes the idea that a mental image could be just an internally generated and inspected copy of the world. (Form a picture of a man climbing a hill with a cane. How do you know what that's a picture of? Could it be a man sliding back down the hill, dragging his cane after him?) Geschwind tells us of the disunity of the mind, how in disturbed patients, the right half always seems to act aggressively towards the left. Again, these problems are all considered in substantial technical detail, made understandable, but not popularized. The specialists manage to enchant us with their recent findings: Brains, like hair color and height, differ from person to person; different internal systems, often at a distance from one another and under different kinds of control, can effect the same behavior, such as facial movement.

A serious problem considered from several angles reflects the eternal mind/body paradox: What is the relation of internal states of the brain to what is represented in the world? How does a symbolic, mental function like thought affect

the physical brain? We aren't quite able to untangle these (probably confused) questions, but these interviews provide a good summary of where we stand. Gregory provides an intriguing analogy, "Brain states represent the world rather as letters on a page represent fiction or truth"—in perception there can be distortion, ambiguity, paradox, and fiction. But how can we will a perception to change, in the manner we will a Necker cube to flip? Gregory resolves this by saying that it's one brain mechanism (will) affecting another part of the brain mechanism (the perceptive process) (p. 60).

Dennett relates how AI's programs, as societies of interacting parts, solve the problem of the homunculus, a way essentially to "get the ideas to think for themselves." The whole activities of the whole system emerge out of the cooperative activities of specialists: "Put together into large armies, mustered like Chinese boxes, ... these stupid elements can exhibit behavior which looks distinctly intelligent, distinctly human." (p. 78) Fodor acknowledges our success with "modularised psychological systems," but is in the end pessimistic of the adequacy of this approach for explaining how the conscious person arises out of this loose federation. How do the modular systems with their highly specialized languages communicate, and what mediates problem solving? Fodor says, "I have no idea how that works, nor am I convinced that the currently available theoretical and experimental techniques are very likely to throw much light on that question." (p. 98)

AI, and information processing in general, is credited, often tacitly, as the wellspring of the "cognitive revolution" (as Bruner puts it). It is interesting to see George Miller using Nilsson's definition of AI ("...intelligent if a person did it..."). Bruner cites the General Problem Solver (Newell & Simon, 1972). Gregory describes how "hill climbing" might be used in perceptual exploration. The computer metaphor is used throughout. But Dennett, interested in modelling human thought, has a disappointed evaluation of the field: "the actual products of AI...are a relatively unimpressive lot... they're typically a bag of tricks... gimmicks and illustrations..." Nevertheless, all of the scientists in their analysis of human thought embrace the methods of AI, nicely described by Fodor: "When we encounter a cognitive system that specialises ..., we can ask quite detailed questions about the kind of cognitive domain it operates in, what's the structure of that domain, what kinds of transformations or operations have to be performed on objects in that domain, what kind of output does the system produce..." (p. 95)

This brings us to the idea of constructive understanding. Miller summarizes it well in his discussion with Bruner: "So we are working towards a psychology which visualises perception as a series of constructions and hypotheses which we act upon until they prove to be erroneous. Then we abandon them, or modify them, and reconstruct them." The ideas are probably best developed in the interview with Gregory, in which examples of illusions and ambiguous figures illustrate the sense in which "perceptions are predictive hypotheses"

and our experience biases our interpretations. Gregory points out that illusions are not a manifestation of the physical structure of our eyes or brain (or video device and computer), rather they arise from inappropriate assumptions incorporated by the strategies or procedures of perception. Importantly, illusions arise from the *power* of the procedures of perception, so any computer program not exhibiting the same or similar misperceptions, would fail to fully share in the advantages of the constructive approach as well.

In most general terms, the idea developed here is that we do not know the world directly, as George Miller puts it, as "responses to penny-in-the-slot stimuli." Rather, perception and understanding—all that we know—go through a conjectural stage of guessing, in which an interpretation is formed, and (perhaps) later stages of matching against experience and revision of the conjecture. Stating again the radical departure from early psychology and behaviorism, George Miller says, "People seldom respond to stimuli. They respond to what they think the stimulus was."

Most important is the feedback operation of matching and revision that occurs. In an important sense, the initial interpretation affects what further information is gained about the world and how it is interpreted. To give a familiar example, a physician's initial diagnoses will affect what further information he seeks about a patient. The revised hypotheses generally follow from the initial guesses, rather than being completely different interpretations.

Initial guesses constitute expectations, which are "matched against the world" for a measure of correctness or veracity. One view is that arousal or shift in attention occurs when an expectation is found to be wrong—in other terms, the model of the world is violated. Bruner develops this from the perspective of information theory, explaining the idea of a "feed-forward" mechanism: "The difference between what was intended and what actually occurred is what we process to correct the movement." This same idea is used by Gombrich to explain how artists study their response to representations (drawings, paintings, etc.) and evolve new techniques to evoke the kind of interpretation they wish to record. Thus, in the simple act of drawing, line and shading are chosen to correct a conjectured (on canvas) representation so that it matches the conjectured (by the mind) representation of what the world is like (or feels like). A drawing and the process of drawing make concrete the notions of mental model and hypothesis refinement that are central to perception, emotional response, high-level problem solving, and social behavior.

Mandler's analysis of emotional arousal fits this model well, though, like most of the discussions in this book, it is somewhat counterintuitive and surprising. For example, a paranoid might have a propensity to view other people as threats. He has constructed a world peopled with rivals, so the smallest action of another appears threatening. With his intentions and reality constantly in conflict, he is continuously "interrupted" by what people do. According to Mandler, this arousal together with the negative evalua-

tion creates the continuous emotional state of rage in the paranoid person. Note that it is the interpretation, based on a fantasy, that produces the emotion, illustrating how one state of mind, a conjectured interpretation, brings about another. The implication sounds like common sense: If you're overly anxious about the world, you probably have a good imagination.

The idea of "constructive understanding" is an application of the hypothesis and test paradigm, which is certainly familiar to AI researchers. However, the breadth of its application by the contributors of this book, as exemplified by Mandler's analysis, is surprising. A summary of how the idea has changed psychology, anthropology, sociology, and psychiatry should make its pervasiveness and importance more evident.

Our early theories of the mind and behavior identified the world with analogies—scientists tended to take the world for what it appeared to be like. In their study of the "representativeness heuristic," Tversky and Kahneman have begun to formalize how our judgment is biased by what is familiar to us (Kahneman, Slovic, & Tversky, 1982). In Miller's interviews, the following common sense analyses are debunked as being based upon a similar erroneous assumption about resemblance. These views are replaced by interpretations founded upon the concept of constructive understanding.

- Anthropologists observed that primitive societies exhibited beliefs about the world that we find in children. Behaving like children, they must be children. Thus, primitives are living in societies that have not grown up, they show us man in his original state.
- Sociologists studying violence at football games adopted the popular viewpoint: These aren't civilized people. They have lost their social mores and regressed to an animal state. These people are exhibiting the former state of man; they are animals.
- You see an artist drawing. He is drawing a landscape. He is trying to copy what he sees. A drawing is a literal copy of the world and drawing is a process of capturing exactly what is in the world.
- A patient with a brain dysfunction is still able to read, though makes a number of errors. Because of the errors, we can infer how she is translating symbols on the page into words. This reveals for us the normal mechanism of reading.
- I have an image in my mind of my office. I can see the desks and windows. It is exactly the way it is in real life. My image is a mental copy of what is in the world. It is an internal picture that I can examine and reason about.
- The guy who shot the President is very confused. He thinks that a movie star loves him. His mind

is not working right. He's feverish; he has an illness, a mental illness. He has a disease.

The contributors to *States of Mind* would argue that each of these interpretations is wrong. What's interesting is that they would each make the same kind of argument. As Geschwind remarks about the problem with analogical interpretation in a narrower context, the great problem with our common sense analyses "is the hidden assumption that similar behaviors always depend on similar mechanisms." (p. 134) Correcting our common sense views, the contributors would argue instead that the behavior or mechanism being observed involves a *constructive process*. The subject is attempting to make sense of the world, and often coping with incomplete information, failure to understand, or lacking an authority by which an interpretation can be made or validated.

Of course, this analysis is replete with reflections upon itself. The contributors are well-aware of the bridge they are forming between common sense reasoning and the scientific process of explanation. They are saying that people, in their ordinary lives as readers, physicians, sports fans, artists, etc. operate within a framework of something like rules or expectations constituting a model of the world. Their behavior follows from this model and critically directs the other information they acquire and how they process it. Formalize the process of recording models and testing them, and you have the scientific method. As Einstein said, "The whole of science is nothing more than a refinement of everyday thinking." And the refinement applies to the content, as well as the process, of thinking.

So how does all of this apply to AI? I'm going to consider some methodological implications, some specific applications, and then broaden the discussion to some of the social and psychiatric issues.

First, with this perspective on the evolution of psychology and the kinds of mistakes made in the past, we might re-examine the information processing model for lingering *non-constructive aspects*. To give one striking retrospective example, consider how the production rule formalism, used for building some "expert systems," was influenced by the work of Newell and Simon. Put a bit simplistically, the production rule formalism has not been completely successful (Clancey, 1983) because Newell and Simon's model of human problem solving is still too close to the stimulus-response view of reasoning. A production system is just a collection of stimulus-response rules with a memory. It doesn't account for the interplay of knowledge organization and reasoning strategy, reasoning about assumptions, reasoning with incomplete knowledge, etc. It might be worthwhile to identify what other aspects of current models of problem solving assume non-constructive processes.

Dennett's discussion of the problems of a "bottom-up" strategy for understanding how the mind works suggests another methodological pitfall. The idea of developing theories of intelligence by proceeding theoretically, by extending logic formalisms, is much discussed today. This

has the potential of developing into something like information theory, which tells us very nicely how to transmit signals clearly and efficiently, but says nothing about “what we might write about, what we might portray, what themes or ideas are going to be presented.” People arguing for an empirical approach to AI are emphasizing the importance of knowing what questions need to be asked and what hypotheses to test. The argument is not against a proper formalism for couching theories, rather against the approach of centering the discussion on the formalism itself.

In some ways, AI’s difficulties might be directly addressed by studies of the nature of thought. For example, one current problem in expert systems research is to develop a single representation that would allow a program both to solve a problem and to explain what it is doing. The inability of people to introspect about underlying cognitive processes involving memory and perception, and indeed Dennett’s suggestion that “it is a constitutive character of having a mind at all,” should make us take pause. There is reason to believe that we need multiple representations, on different levels of abstraction (e.g., as used by Brown and Burton in SOPHIE, to relate a numeric, FORTRAN-coded simulation of an electronic circuit to laws and component models (Brown, 1977)). Or, perhaps more accurately, the description is generated by running the problem-solving procedure, making observations, and conjecturing about what is happening. So maybe introspective explanation of reasoning is a constructive process, too.

To paraphrase Bruner, “The logic of problem solving is not a process, it is the *characterization of the results* of processes at work.” (p. 35) Relating waves to procedures and particles to descriptions, the apparent paradox is reminiscent of Heisenberg’s result: “Certain kinds of knowledge, canonically speaking, cannot be considered simultaneously.” Put another way, reflecting on a procedure and running it cannot go on at the same time—you’d trip over your feet. The implications are surely not worked out yet, but as Dennett points out, AI’s goals place us squarely in the camp of Hume, Descartes, and Kant, wrestling with the old questions of consciousness and representation.

We can be more specific about some implications for research in the area of learning. Today we can partially answer Plato’s “Paradox of Meno”—how can you learn something you don’t know? Our answer is that you formulate a guess and fix it. You don’t *recall* it full-blown; you aren’t *given* it full-blown; you *construct it from cues*. One implication is that, to be successful (and efficient), explorative learning requires a model for guiding search—some kind of idea of the form of the solution. You need to know just what you’re supposed to gain from the experience. You benefit from having some way of “framing” the cues. Bruner relates an interesting experiment in which infants appear to be taking cues from where adults are looking. Other experiments suggest that infants won’t explore complex environments for which they lack perceptual models.

Gombrich’s analysis of drawing suggests an intriguing

analogy with learning. He describes how an “artful blurring . . . invites the beholder to increase his contribution.” The beholder’s share is chosen by the artist in a way reminiscent of a teacher framing a lesson so that the student will realize the point for himself. The art is in knowing what distinctive features should not be left out and how to point to the intended whole. In this way, artists and teachers *manipulate* the process of constructing an understanding. In recent research on learning, Van Lehn has called these directive, essential features that contribute to meaningful encounters “felicity conditions,” drawing parallels with the “conversational postulates” of Grice (VanLehn, 1983). It is the strength of Miller’s book that it is conducive to conjecturing and studying analogies between areas as apparently diverse as drawing and teaching.

I have to this point drawn most of my examples and discussion from the mental and perceptive section of the book (George Miller, Bruner, Dennett, Gregory, Fodor, Geschwind, Gombrich). To conclude, I want to develop the social and psychiatric ideas a bit and relate them to the more long-term problems of constructing machines that think.

The discussions of psychiatry by Hampshire, Farrell, Segal and Szasz, coming at the end of book, lead one to wonder how psychiatry fits into the evolving picture of “man, the maker of meaning.” The problem is that Freud’s analysis tries to bridge poetry, physiology, and physics, unifying myths with instinct and forces. Hampshire summarizes it well: Freud “is in an interstice—a kind of gap between purely imaginative enquiries and scientific ones. [His theories] . . . may have a transforming effect on the way we look at people, and they make a claim to be true, but true in a sense which doesn’t reduce to a set of proven evidenced propositions.” (p. 114) In the reformulation of Freudian theories by people like Klein, we find an analysis that nicely parallels our modern understanding of perception and problem solving.

The important step is to view psychological problems as revealing a problem-solving process, specifically an attempt to make sense of conflicting intentions, behavior, and experiences. The individual’s problem is usefully addressed at the level of meaning (or the inability to construct it), not in terms of energy or “libidinal forces.” Rather than responding to instinct, the infant, for example, is trying to deal with a given situation. Confronted with threatening interpretations and conflicting desires, he experiences emotional tension and discomfort. This process of meaning construction is further engaged in by the psychiatrist who constructs an explanatory analysis, and again on a larger scale by the populace, who can weave personal stories following the tenets of the “Freudian church.” In this view, the main contribution of traditional psychoanalysis is that it reduces the tension of being confused; it provides a framework that satisfies our need to understand ourselves and integrate our experiences. And as Farrell says, “*Conviction* of the truth is important, not whether the construction is really true or not.”

Going beyond a superficial summary, the main points of the contributors are:

- In the face of contrary needs and confused interpretations, we need some way to achieve integration.(Segal)
- In the face of violated expectations or unsatisfied desires, we need to reconcile our fantasies and plans with reality.(Segal)
- In the face of failure to be effective, we need some way to sustain belief that the world can be understood and that our knowledge and actions will be effective. (Geertz)
- In the face of ultimate questions of authority and meaning, our existential difficulties, we need a belief framework in which rationality can continue to exist and by which we can endure uncertainty.(Geertz)

It is beyond the scope of this review to provide a full discussion of these points. The contributors provide numerous examples drawing from diverse sources:

- theories of primitive man—Geertz: “They are contemporaries of ours, not our ancestors.”(p. 198) “The ritual activity is not conceived as instrumental. . . it’s an attempt to display to themselves and to reinforce in themselves the fact that the world is the way they think it is.” (p. 204)
- infant development—Segal: “The mind is an achievement; doing something which is restorative means being in touch with the psychic pain and doing something about it.”(p. 263)
- social psychology—Harré: “Handshakes are psychologically effective in so far as they are interpreted. . . (they) create a fragment of social order.” (p. 158) “A skilled social talker can recreate the past.” (p. 161)
- the ethics of psychiatry—Szasz: “The psychiatrist, for his part, does two things. He takes innocent persons and deprives them of their liberty; that’s called civil commitment. And he takes guilty persons and claims that they are innocent; that’s called the insanity defense.” (p. 276) “What does the patient want to do? . . . The decision about whom he turns to for help must grow out of who he is.”(p. 289)

The common methodological approach is well-stated by Geertz: “Try to understand what the people themselves think they are up to and try to explain *that*, try to get some sense of what is going on among *them*.” (p. 205)

The overarching lesson for AI is that these moral difficulties are not the result of some lingering animal nature in man, some kind of irrationality that we are cursed with, tucked away in some reptilian inner brain. That kind of popularized, “resemblance-is-the-fact,” analysis ignores the essential dilemmas of man as an independent maker of meaning: our inability to know fully how the world functions,

why we exist, and why we behave the way we do. Any program that we construct with the self-awareness to examine its own behavior, gather information and solve problems, and attempt to learn from experience must confront these same difficulties. Indeed, as suggested by perceptual illusion, we might say that the ultimate test of scope of intelligent activity is being confused about one’s own meaning and purposes.

The best summary is again provided by Geertz: “. . . Life for the acculturated animal, always involves at the center of it, a search for at least some level of coherence of meaning, for the significance of what’s going on. A sense that we know who we are, that we know what the world is *like*, and that we have some purchase on what’s going on; that it isn’t one damn thing after another.” (p. 208) But Szasz insists that we call rituals for what they are: “Life is something we must endure. There is no solution for it.”(p. 290)

AI programs cannot avoid these problems. Partially tested, conjectured models of the world and fantasies (plans, expectations) make up the “irrational” part of what we know. Just like people, programs will need “a capacity for reality testing and distinguishing the psychic from the external,” for example, by “lessening the concreteness” of simplistic analogies. Programs acting as independent agents in networks must cope with conflicting goals and limited resources. They will have to decide when and whether to lie. Programs will have to deal with matters of pride and responsibility because these are naturally part of coping with a framework of rules, rights, and duties (Harré).

Weizenbaum tells us that we shouldn’t give programs social responsibility because “theirs must always be an intelligence *alien* to genuine human problems and concerns.” (Weizenbaum, 1976; p. 213) Yet, he acknowledges that the intelligence of man is constructed, arising from the problems he confronts. In the face of the analysis given here, we might suppose that the biological and emotional differences between man and machine, which Weizenbaum focuses upon, might pale in comparison to the shared difficulties of maintaining rationality in a given society.

If we treat programs as if they were responsible, they will be caught “like dogs and children” in trying to live up to moral expectations. As Harré says, “Everything that moves, be it mechanical or animal, gets sucked in.” This follows from the observation, to paraphrase Harré, that the “mind comes into being by virtue of holding theories about itself.” We don’t (and can’t) supply morality full-blown to a child or a program; the reflective, acculturated problem solver is led to consider and formulate his own model of the moral expectations others hold about him. As Weizenbaum says, “. . . his humanity . . . depends crucially on his seeing himself, and on his being seen by other human beings, as a human being.”

A program caught trying to satisfy multiple goals with limited knowledge and resources will naturally become embroiled in moral difficulties, and that independent struggling for meaning fits in large part our view of what an individual

mind is. Faced with similar social difficulties, the mind of a computer on earth might not be as alien as Weizenbaum suggests: Moral differences might not be as extreme as between human cultures themselves. But what is most intriguing, by not giving computers responsibility to act morally, as Weizenbaum so strongly recommends, we guarantee that they will never understand or be sympathetic to our most difficult problems. This severely bounds, but perhaps not impractically, the help they might provide as assistants in the office, classroom, and home.

To step down off this rather lofty plane, I will conclude with some more mundane observations about Jonathan Miller's book. As can be seen, the book is a good source of historical information, well-suited for a seminar or a companion to a cognitive science survey course. Miller's success stems from having done his reading first, and so he was prepared to draw out his contributors to help them make their best points. The strength of his broad background as physician, medical historian, playwright, and director clearly shows.

The most insightful, balanced remarks are made by Harré and Geertz. Szasz's extremely strong, persistent views could hardly be characterized as balanced, but the strength of his convictions and originality make this interview (on the "myth of mental illness") worth reading again and again. But there's a lot of tension in the conversation with Hinde, whose research borders most closely upon traditional psychology. Miller seems disappointed with how his former teacher has changed, and he perhaps should have left this one out.

The humanism of these men seems to have evolved from their studies of human nature. Miller points out to Geertz that he appears to be bending over backwards to avoid being condescending; these men have the wisdom to turn their own analysis upon themselves. Perceiving the search for "honor" as the central recurring theme of man's experience, they are careful to approach other cultures and scientific disciplines in terms of their perceived problems and what they were trying to do. This could explain the admirably balanced view of traditional psychology that appears throughout the book. Harré states the moral: "If there is any distinction between 'primitive' thought and 'modern,' it lies with the possibilities for self-criticism of thought made possible by contact with other ways of looking at things." (p. 202) In fact, this book is really a testimony to self-knowledge, and it has that persuasive kind of honesty about it.

—William J. Clancey
Department of Computer Science
Stanford University

References

Brown, J. S. (1977) Remarks on building expert systems. In "Reports of panel on Applications of Artificial Intelligence," *Proceedings IJCAI 5*, 994-1005.

- Clancey, W. J. (1983) The epistemology of a rule-based expert system: A framework for explanation. *Journal of Artificial Intelligence*, 20(3), 215-251.
- Kahneman, D., Slovic, P., & Tversky, A., (Eds.) (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.
- Newell, A. & Simon, H. (1972) *Human Problem Solving*. Englewood Cliffs: Prentice-Hall, Inc.
- VanLehn, K. (1983). Human procedural skill acquisition: Theory, model, and psychological validation. *Proceedings of the National Conference on Artificial Intelligence*, 420-423.
- Weizenbaum, J. (1976). *Computer power and human reason*. San Francisco: W. H. Freeman and Company.