Epilogue

We began this book with the question: *To what extent is it possible to model intelligence as an information processing activity that can be carried out by a machine?* This epilogue restates and summarizes the most important views and arguments relevant to this *computational hypothesis*, and examines whether it is possible to construct an "intelligent machine" that can function in the world. A discussion of the possibility of a robotic entity is appropriate because such a device would have to draw on the entire range of intelligent capabilities, from reasoning to language to vision, and the problem of integrating these capabilities brings to the fore many of the deep questions that arose in the individual chapters.

The human brain, a three-pound lump of biological tissue, is perhaps the most amazing object in the universe—it sees, hears, thinks, creates, and even seeks an explanation for its own existence. Can it really be the case that in spite of its remarkable properties the brain is simply a "machine" that carries out complex computational procedures? This question has been the basis of the intellectual quest that we have pursued in this book. An affirmative answer seems to degrade the concept of humanity to a point unacceptable to most people despite the fact that we are generally willing to accept the idea that other living creatures may indeed be very sophisticated machines. If we refuse to accept the computational hypothesis, then only mysticism is left, a situation unacceptable to most scientists.

PHILOSOPHICAL AND CONCEPTUAL QUESTIONS

In this epilogue we deal with major questions raised at various points in the book, some of which reflect a basic concern about the brain as a machine, some that concern the machine as it affects society, and some about the field of artificial intelligence in general. For each question we will first summarize our beliefs and the points of view of others in the field, and then indicate the impact of the question on the possibility of creating a humanlike autonomous robotic device.

Because we are dealing with ultimate limits of an intelligent machine, we quickly find ourselves delving into rather deep philosophical questions that concern such topics as understanding, consciousness, language, and mind. These questions arise when the behavior of a machine causes us to ask, "Does the *behavior* of the machine mean that the machine is *really* capable of X," where X is understanding, consciousness, thinking, etc. Some of the questions that have been discussed for thousands of years have taken on a new interest and importance due to artificial intelligence (AI).

True Understanding

Can a machine truly understand?

One point of view concerning machine understanding is stated by Simon [Simon 81]: "At the root of intelligence are symbols, with their denotative power and their susceptibility to manipulation." Such symbol manipulation can be carried out by a brain or by various forms of computers. Simon therefore sees no reason why machines should not be capable of human forms of understanding the world. In Chapter 6 we discussed the strongest attack on this point of view, Searle's Chinese room metaphor [Searle 84], which asks whether one can ascribe "understanding" to mere symbol manipulation. An important point of the Chinese room metaphor is to show that when we limit a person's role to strict symbol manipulation, this activity adds nothing to the person's ability to truly understand the task at hand. Although Searle makes no attempt to supply an operational definition of understanding (or even a set of necessary conditions), what makes the metaphor so powerful is that we ourselves would not require or attain an understanding of Chinese if placed in this situation. If we, as conscious thinking creatures do not exhibit understanding in the presence of such formal symbol manipulation, how can a machine?

Even though he does not make the explicit point, Searle appeals to the idea that consciousness plays a prime role in human understanding. (In a slightly modified way, this is Descartes' mind-body problem again; see Chapters 1 and 2.) Consciousness seems to be some coherent *whole* that cannot be decomposed into electronic switches, or computer instructions, or even subsystems, and thus is not subject to the type of reductionist analysis often used in science. Why do we have the introspective ability called "consciousness?" Minsky [Minsky 68] believes that the evolutionary role of consciousness has been to give access to otherwise inaccessible modules, so as to debug, reprogram, or retrain them. He suggests that an organism would be better able to survive if it has a model of itself. Dennett [Dennett 84] imagines the development of consciousness as follows. People found that they could get advice on future actions by asking other people. At some point it was found that even if nobody is around, a person's problem solving activity is aided by asking questions out loud. Gradually, it was realized that the same effect can be achieved by internally asking the questions,

> Under what conditions would the activity of asking oneself questions be useful? . . . Crudely put, pushing some information through one's ears and auditory system may stimulate just the sorts of connections one is seeking, may trip just the right associative mechanisms, tease just the right mental morsel to the tip of one's tongue. One can then say it, hear oneself say it, and thus get the answer one was hoping for.

One additional function that consciousness appears to serve is that of a high-level executive, or even an external observer, who monitors the performance of the body in relation to its high-level goals, and makes appropriate suggestions, e.g., as in the case of a runner verbally "telling himself" to go faster (See Box 2-3).

Sometimes if we know the structure of a machine or the background of a person we can make a decision concerning the nature of the understanding involved. For example, suppose we know that we are dealing with a person who has been blind from birth and has never experienced color. He has been provided with an instrument that senses color and communicates the results in Braille. The person can point the sensor at objects, and can then report color as well as a sighted person. No matter how good his operational performance is, we still feel that he has no understanding of what it means to see the world in color.

Sloman [Sloman 85] would not agree with the above point of view. He argues that there is no clear boundary between things that do, and things that do not understand symbols. He feels that our ordinary concept of "understanding" denotes a complex cluster of capabilities, and different subsets of these may be exhibited in different people, animals or machines; to ask "Which are necessary for *real* understanding?" is to attribute spurious precision to a concept of ordinary language. Thus, in the example above, Sloman would say that a congenitally blind person may attach meanings to color words not too different than a sighted person, because much of the meaning resides in the rich interconnection with concepts shared by both.

The potential role of "consciousness" in a robotic device is uncertain. Certainly, it is possible for each subsystem of a robot to signal its status and goals to other subsystems and to attain the external appearance of conscious behavior, but what have we lost by implementing the *appearance*, but not the *substance*, of conscious awareness? For example, is it imaginable that such an implementation would be capable of autonomously conceiving of, and printing, the message "I think (compute?), therefore I am"?

The issue of understanding is a crucial one in the robotic domain, since it brings to the fore the basic question as to whether "true understanding" is required to achieve high levels of performance, i.e., to give explicit shape to high-level goals in an uncertain environment, and whether robotic devices are responsible enough to take their place in society, as discussed in the following paragraphs.

Societal Aspects

Can a machine take a role in society as if it were a person?

Suppose we reach a point in time when a robotic entity can be placed in a position to carry out decision making in government, banking, medicine, law, or the military. Could the robot perform as if it were a person. No! answers Weizenbaum [Weizenbaum 76], who points out that knowledge about a society "is acquired with the mother's milk and through the whole process of socialization that is itself so intimately tied to the individual's acquisition of his mother tongue. It cannot be explicated in any form but life itself." Thus, an American judge, no matter what his intelligence and fairmindedness, could not sit in a Japanese family court. His intelligence is simply *alien* to the problems that arise in Japanese culture. "Whatever intelligence a computer can muster, however it may be acquired, it must always and necessarily be absolutely alien to any and all concerns." This question is not a theoretical one, for we already have medical and business decision making programs in use that affect human health and welfare.

Weizenbaum and many others believe that true understanding can only be gained by actually experiencing the world and thereby developing an internal database that represents these experiences. A child builds up knowledge of the world by exploring and learning. Similarly, it is necessary for a robotic device to build up its database of world knowledge by learning about its environment, since manual entry of sufficient relevant information about the world is impractical. While present-day machine learning consists merely of determining the specific parameters and relationships for a representation that has been chosen by the designer of the machine, we have no reason to suspect that machines are inherently incapable of the more powerful forms of learning needed to counter Weizenbaum's criticisms.

If we change the question in this subsection to, *Can a machine play a useful role in society?* rather than requiring the machine to function as a person, then the strict requirements for deep understanding of the society indicated by Weizenbaum can be relaxed. The designer has a much simpler task, and quite useful devices can be constructed. Weizenbaum would still feel that society must carefully control the use of such devices, and avoid applications where the consumer might ascribe more intelligence to the device than is warranted.

The Frame Problem

How can a computer focus its attention on just those aspects of a given problem that are relevant to its solution?

This is the *frame problem*: How can an intelligent system, in trying to solve a problem or carry out an action, know what information in its database should be ignored, and what is relevant to the problem at hand? The question, "What was Benjamin Franklin's telephone number?," given in Chapter 3, is a relevant example. Somehow people know that they must invoke the relevant information of Franklin's date of birth and the date of the invention of the telephone to solve the problem.

To fully represent the conditions for the successful performance of an action, an impractical and implausible number of qualifications would be required. McCarthy [McCarthy 80] gives as an example the many qualifications that would have to be supplied for the simple act of a boat crossing a river—the oars and oarlocks must be present and unbroken, the boat cannot have a hole in it, cannot be filled with rocks, etc. Since many other qualifications can be added, the rules for using a rowboat become almost impossible to apply unless some focusing mechanism is available.

The frame problem is of crucial importance to an autonomous robot. It must "think before it leaps," coming up with reliable but not necessarily foolproof expectations of the effects of its actions. This process must be carried out in an acceptable amount of time, taking into account everything in its database that is relevant to the proposed action. The big problem is knowing what is indeed relevant, since the time constraints will not permit examination of the implication of every fact in the database. Many AI systems have the unfortunate characteristic that increasing the amount of knowledge in their database degrades, rather than improves, their performance.

There are currently several approaches that try to deal with the frame problem. One scheme uses the attention-focusing ability of stereotypes, such as the Schank script approach described in Chapter 6. A script attempts to define and organize all of the relevant information for a given situation, but there is still the problem of knowing which of many possible scripts is relevant to a complex situation. McCarthy has proposed [McCarthy 80] *circumscription*, a form of nonmonotonic reasoning augmenting ordinary first-order logic as a mechanism that can be used for "jumping to appropriate conclusions." Circumscription conjectures that the only entities that can prevent an action are those whose existence follows from the facts at hand. Loosely speaking, it is a "don't-go-looking-for-trouble" approach. In the example of the boat, if no lack of oars or other circumstance preventing the use of the boat has already been deduced, then conclude that the boat is usable. A program must contain heuristics for deciding what circumscriptions to make and when to withdraw them.

It may be that the frame problem is an artifact arising from the symbolic representations used in AI. Such representations eliminate the implicit connections that exist between objects in the world, and much effort is required to "put Humpty Dumpty back together again." In particular, the logic formalism does not capture the implicit relationships between things in the world nor permit the type of approximation that seems to be required for everyday reasoning [Pentland 83]. The seriousness of the frame problem is stated by Dennett [Dennett 84a]:

> It appears at first to be at best an annoying technical embarrassment in robotics, or merely a curious puzzle for the bemusement of people working in AI. I think, on the contrary, that it is a new, deep epistemological problem—accessible in principle but unnoticed by generations of philosophers—brought to light by the novel methods of AI, and still far from being solved.

The frame problem is closely related to the decomposability question discussed below.

Decomposing Intelligence

Is intelligent behavior decomposable?

As indicated in Chapter 3, most formal analysis begins by decomposing a problem into manageable portions (see also the discussion of the partitioning problem in Chapters 8 and 9 and the necessity for the independence assumption in both formal and probabilistic reasoning, Chapter 4). We often assume that the whole is made up of its parts and that such decomposition can be carried out for even the most complex of situations. But suppose, in fact, that much of intelligent behavior is not decomposable, and that in partitioning behavior for mechanization purposes, one is incorrectly making an assumption of independence of the parts. Until we are able to deal with the entire unpartitioned problem, we would not be able to achieve human performance for such tasks. Dreyfus [Dreyfus 79] has stated,

> . . . . all aspects of human thought, including nonformal aspects like moods, sensory-motor skills, and long-range self-interpretations, are so interrelated that one cannot substitute an abstractable web of explicit beliefs for the whole cloth of our concrete everyday practices [p.54] . . . . Since intelligence must be situated it cannot be separated from the rest of human life [p.62] . . . . If one thinks of the importance of sensory-motor skills in the development of our ability to recognize and cope with objects, or of the role of needs and desires in structuring all social situations, or finally of the whole cultural background . . . . the idea that we can simply ignore this know-how while formalizing intellectual understanding as a complex system of facts and rules is highly implausible [p63].

In a more intuitive sense, the naturalist John Muir noted that everything in the universe is attached to everything else. If it turns out that it is not possible to decompose many of the processes required for an entity to function in the world, then non-learning robotic devices will be limited as to the type of reasoning and problem solving they can carry out (because of practical limits on what can be designed into them). In particular, if intelligence is not decomposable, it is not incrementally achievable: the idea of building a partially intelligent robot would make as much sense as digging half a hole.

Again we must stress that we are talking about highly intelligent entities. Once we relax that criterion, then it is surely possible to develop useful devices based on reductionist principles.

Language and Thought

What is the relationship between language and thought?

This question has been discussed for thousands of years, and the final answers have still not been found. In many myths about the creation of man, language is taken to be one of man's inherent characteristics like vision or hearing.[1] Some of the many theories about the relation of language and thought include:

1. *Thinking is a form of subvocal speech.* John B. Watson, the father of the psychological school known as behaviorism, believed that all thinking was a form of motor behavior that had been conditioned to stimuli. When challenged that no motor behavior can be seen while thinking, he claimed that thinking was subvocal speech i.e., motor behavior that could not be noticed. Thus, people who were engaged in thinking activities were just softly talking to themselves [Watson 30]. This point of view was discredited in 1947 by an experiment in which a subject was able to think despite complete paralysis of his musculature [Smith 47].
2. *Language influences a person's thinking and his perception of the world,* the Whorfian hypothesis [Whorf 56]. This hypothesis was discussed in Chapter 6. No definitive experiment has proved or disproved this hypothesis because it is difficult to design an experiment that is not subject to the many cultural and environmental factors that bear upon language. Thus, each experiment that supposedly proves or disproves Whorf's hypothesis can be attacked by opponents who cite confounding factors such as the environment in which the subjects grew up, possibility of misunderstandings of the experiment by the subjects, or similar performance by an entirely different type of subject.
3. *Language is shaped to fit the prelinguistic thoughts that are to be communicated,* a point of view argued as long ago as Aristotle. In the seventeenth century John Locke wrote that man's organs were fashioned to "form articulate sounds" and he was given the ability to "use these sounds as signs of internal conceptions and make them stand as marks for the ideas within his own mind." This view says that thinking comes before speech, and that our language is shaped by the thoughts that we want to convey. Studies that support this theory attempt to find a structural commonality in language structure among the world's languages that would indicate that language is shaped by thought. For example, Greenberg [Greenberg 63] showed that for 98 percent of the world's languages the subject precedes the object—an indication that in constructing linguistic structures, we use a word order that first establishes what the sentence is about.
4. *Language mechanisms of the brain are unlike those of any other cognitive skill,* has been an important theme in Chomsky's work. This theory says that language, like vision, uses genetically determined special neural structures [Caplan 82, Lenneberg 67]. Since language is species-specific to man, it is not possible, as in vision, to carry out animal experiments to gain insight as to the relation between linguistic stimulation and neural structure. Therefore, the advocates of this position must reason from the following basic facts: (1) children rapidly learn language by being immersed in a linguistic environment, (2) modern linguistics has shown how difficult sentence comprehension really is, and (3) children will learn the particular

[1]The history of the biological basis of language by Otto Marx, given in Appendix B of E. H. Lenneberg's *Biological Foundations of Language* [Lenneberg 67] makes fascinating reading.

language of their environment. Thus advocates postulate special brain structures for language that somehow have the power to deal with the complexity of language, and yet have the *plasticity* so that there is no bias toward any one particular language. Opponents feel that learning a language is a special case of very general cognitive ability, and that language learning is merely the application of general learning procedures to the special case of learning to talk [Anderson 80]. At the present time the conflict is not resolved. As Marshall states: "We just do not know whether the neurons, synapses, transmitter substances, patterns of connectivity, and so forth in the language area of the brain differ in important respects from those characteristic of other parts [Marshall 80]."

It is interesting that most of the discussions about thought and language give short shrift to thought and image. If it turns out to be the case that thinking involves languagelike processes, as suggested by Luria [Luria 73], then our present-day approach to robotic reasoning based on symbol manipulation would be justified. Yet, there is the nagging thought that animals, who do not have the language facility of humans, are obviously able to reason, survive, and to achieve goals. Can it be that they are using image-based reasoning? A discussion of the "representation of thought" question is continued below.

Representation in the Brain

What is the nature of representation used by the brain?

Fodor [cited in Miller 83] states: "I suspect that the representational system with which we think, if that's the right way to describe it, is so rich that if you think up any form of symbolism at all, it probably plays some role in thinking." The mind apparently uses two major representations, propositions and images (or at least representations that are isomorphic to images) (see Chapter 1). These two representations have been used in AI, and there are strong advocates of each approach. For example, Nilsson [Nilsson 83] asserts that AI is the study of how to acquire and represent knowledge within a logiclike propositional formalism, and the study of how to manipulate this knowledge by use of logical operations and rules of inference; he does not see a need for additional (iconic) representations. Some feel that this point of view is extreme. Pentland and Fischler [Pentland 83] stress that multiple representations are necessary. Weaknesses in the propositional formalism can sometimes be eliminated through the use of an auxiliary isomorphic representation, since a representation whose structure "mirrors" (is *isomorphic* to) some properties of the domain being represented is able to *implicitly* represent those properties preserved by the isomorphism. "How is it that an isomorphism-based reasoning process can succeed where the theorem-proving approach fails? The trick is that the isomorphic approach makes use of the *semantics* of the problem domain 'built into' the representation to express useful *approximations*, . . . while logical systems admit no such *approximations* [Pentland 83]." However, Fodor [cited in Miller 83] has observed that one problem about imagery is its limited "expressive capacity," in comparison with language: "There is, for example, no difficulty in *saying* of a man that he is not scratching his nose, however difficult it may be to have an *image* of his not doing so."

The *connectionist* approach of Feldman [Feldman 82] is another possible computational alternative to logical reductionism. The fundamental premise of connectionism is that individual neurons do not transmit large amounts of symbolic information. Intead, they compute by being appropriately connected to large numbers of similar units. This approach employs a representation for expressing local constraints and uses a parallel computational process based on *relaxation* (see Appendix 9-1) to achieve some overall goal.

A robotic device will certainly use both propositions and iconic representations, since each serves best to solve distinct types of problems. However, as discussed in Chapters 8 and 9, progress has been slow in understanding how to effectively represent iconic knowledge in the machine.

Future Prospects

What are the future prospects for AI?

We are a long way from having an integrated robotic system that can function in the real world, even at the level of a five-year-old child. Depending on one's level of optimism or pessimism, this can be viewed either as offering much exciting future research potential, or of indicating that "we may never get there from here." A recent paper brought together comments from various experts in the field of AI [Bobrow 85]. Some of these are given below:

> *Bernard Meltzer*: With very few exceptions, all of AI until now has been concerned with what Freud termed secondary processes of the mind, that is, those concerned with logical, rational, reflexive or potentially reflexive, commonsense thinking; it has neglected the primary processes, that is those concerned with apparently non-rational, non-reflexive thinking that results for instance in new metaphors, shafts of wit, jokes, dreams, poems, brain-waves, neuroses and psychoses.
>
> *Nils Nilsson*: [I predict] better understanding of the relationships between perception and reasoning, codification of a large and useful store of commonsense knowledge, significant progress on such conundrums as the frame problem and nonmonotonic reasoning, and large-scale systems based on the "belief-desire-intention" model of intelligent agents.
>
> *Terry Winograd*: There are two quite different starting points to define AI—the dream and the technology. As a dream, there is a unified (if ill-defined) goal for duplicating human intelligence in its entirety. As a technology, there is a fairly coherent body of techniques that distinguish the field from others in computer science. In the end, this technology base will continue to be a unified area of study with its special methodology. We will recognize that it is not coextensive with the dream, but it is only one (possibly small) piece. . . .

A SUMMARY

We have completed our intellectual journey. Some might feel that they only retain snapshots as a result of their travels—and that they bring back no coherent story. We could have attempted to provide such a story by organizing all of intelligent behavior under a single theory as the "logic imperialists" have suggested [Nilsson 83, Pentland 83]. Or we might have presented a collection of subtheories, such as those of Marr [Marr 82] in vision, or Chomsky in language [Chomsky 75]. Alas, none of these more encompassing theories are believable in the light of our current knowledge. One might be tempted to use the robot as an integrating framework, but here, again, there is no encompassing theory: robotics at present is an application domain for the computational techniques we have presented, rather than a primary source of intellectual ideas. Therefore in this final chapter we tried to weave together the various separate threads that were examined earlier in the book, using major intellectual questions as a focusing mechanism.

Finally, it should be emphasized that in this epilogue we were mostly dealing with questions of the *ultimate* role of the computer as an intelligent device. Thus there may have been a tone of pessimism in some of the views presented here. However, in the near term, for limited domains and less ambitious goals, we are confident that efforts to achieve these goals will result in advances that will both contribute to and quite likely revolutionize society, even if humanlike robots are not created.

Bibliography

[Alford 78] D. K. H. Alford. The demise of the Whorfian hypothesis. In *Berkeley Linguistic Society Proceedings*, pages 485–499. Berkeley Linguistic Society, Berkeley, Calif., 1978., volume 4.

[Allport 80] D. A. Allport. Patterns and actions: Cognitive mechanisms are content-specific. In G. L. Claxton (editor), *Cognitive Psychology: New Directions*. Routledge & Kegan, London, 1980.

[Amarel 68] S. Amarel. On representations of problems of reasoning about actions. In D. Michie (editor), *Machine Intelligence 3*. American Elsevier, New York, 1968.

[Anderson 81] J. R. Anderson (editor). *Cognitive Skills and Their Acquisition*. Lawrence Erlbaum Associates, Hillsdale, N.J., 1981.

[Anderson 83] J. R. Anderson. *The Architecture of Cognition*. Harvard Press, Cambridge, Mass., 1983.

[Anderson 85] J. R. Anderson. *Cognitive Psychology and Its Implications*. W. H. Freeman., San Francisco, 1985, 2nd edition.

[Andrews 70] H. C . Andrews. *Computer Techniques in Image Processing*. Academic Press, New York, 1970.

[Appelt 85] D. Appelt. Planning English referring expressions. *Artificial Intelligence* 26:1–34, 1985.

[Arbib 72] M. A. Arbib. *The Metaphorical Brain*. Wiley, New York, 1972.

[Arnheim 69] R. Arnheim. *Visual Thinking*. University of California Press, Berkeley, Calif., 1969.

[AI 84] Artificial Intelligence. Special Volume on Qualitative Reasoning about Physical Systems. *Artificial Intelligence*. 24, 1984.

[Asher 66] H. Asher. *Experiments in Seeing*. Fawcett Publications, Greenwich, Conn, 1966.

[Aslin 83] R. N. Aslin, D. P. Pisoni, and P. W. Jusczyk. Auditory development and speech perception in infancy. In M. M. Haith and J. J. Campos (editors), *Handbook of Child Psychology: Infancy and Developmental Psychobiology*. Wiley, New York, 1983.

[Atkinson 68] R. L. Atkinson and R. M. Shifrin, Human memory: proposed system and its control process. In K. W. Spence and J. T. Spence (editor), *The Psychology of Learning and Motivation: Advances in Research and Theory*. Academic Press, New York, 1968.

[Ballard 81] D. H. Ballard. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition* 13:111–122, 1981.

[Ballard 82] D. H. Ballard and C. M. Brown. *Computer Vision*. Prentice-Hall, Englewood Cliffs, N.J., 1982.

[Banks 83] M. S. Banks and P. Salapatek. Infant visual perception. In M.M. Haith and J. J. Campos (editor), *Handbook of Child Psychology: Infancy and Development Biology*. Wiley, New York, 1983.

[Barlow 53] H. B. Barlow. Summation and inhibition in the frog's retina. *Journal of Philosophy* (London) 119:69–88, 1953.

[Barrett 79] W. Barrett. *The Illusion of Technique*. Anchor Books, New York, 1979.

[Barrow 78] H. G. Barrow and J. M. Tenenbaum. Recovering intrinsic scene characteristics from images. In A. R. Hanson and E. M. Riseman (editors), *Computer Vision Systems*. Academic Press, New York, 1978.

[Barrow 81] H. G. Barrow and J. M. Tenenbaum. Computational vision. *Proceedings of the IEEE* 69:572–595, 1981.

[Beck 83] J. Beck, B. Hope, and A. Rosenfeld (editors). *Human and Machine Vision*. Academic Press, New York, 1983.

[Begbie 73] G. H. Begbie. *Seeing and the Eye*. Anchor/Doubleday, New York, 1973.

[Binford 81] T. O. Binford. Inferring surfaces from images. *Artificial Intelligence* 17:205–245, 1981.

[Biological 68] Biological Sciences Curriculum Study Group (C. A. Welch, D. I. Arnon, H. Cochran, F. C. Erk, J. Fishleder, W. V. Mayer, M. Pius, J. Shaver, and F. W. Smith, Jr.). *Biological Science: Molecules to Man* Houghton Mifflin Co., Boston, 1968.

[Block 81] N. Block (editor). *Imagery*. MIT Press, Bradford Books, Cambridge, Mass., 1981.

[Bobrow 68] D. G. Bobrow. Natural Language Input for a Computer Problem-Solving System. In M. Minsky (editor) *Semantic Information Processing*. MIT Press, Cambridge, Mass., 1968, Chapter 3.

[Bobrow 85] D. G. Bobrow and P. J. Hayes. Artificial intelligence—where are we? *Artificial Intelligence* 25:375–415, 1985.

[Boden 77] M. A. Boden. *Artificial Intelligence and Natural Man*. Basic Books, New York, 1977.

[Boden 81] M. A. Boden. *Minds and Mechanisms*. Cornell University Press, Ithaca, N.Y., 1981, chapter 11.

[Boolos 80] G. S. Boolos and R. C. Jeffrey. *Computability and Logic*. Cambridge University Press, Cambridge, England, 1980. 2nd edition.

[Brachman 85] R. J. Brachman and H. J. Levesque (editors). *Readings in Knowledge Representation*. Morgan and Kaufmann Publishers, Los Altos, Calif., 1985.

[Brady 82] M. Brady. Computational approaches to image understanding. *Computing Surveys* 14:3–71, March 1982.

[Brooks 81] R. A. Brooks. Symbolic reasoning among 3-D models and images. *Artificial Intelligence* 17:205–244, 1981.

[Brown 58] R. Brown. *Words and Things*. The Free Press, New York, 1958. chapter 7.

[Brown 77] R. Brown. *A First Language: The Early Stages*. Harvard University Press, Cambridge, Mass., 1977.

[Brown 80] J. W. Brown. Language, brain, and the problem of localization. In D. Caplan (editor), *Biological Studies of Mental Process*. MIT Press, Cambridge, Mass., 1980.

[Brownston 85] L. Brownston, R. Farrell, E. Kant, and N. Martin. *Programming Expert Systems in OPS5*. Addison-Wesley, Reading, Mass., 1985.

[Buchanan 78] B.G. Buchanan and E.A. Feigenbaum. DENDRAL and METADENDRAL—Their applications dimensions. *Artificial Intelligence* 11:5–24, 1978.

[Buchanan 84] B. G. Buchanan and E. H. Shortliffe. *Rule-Based Expert Systems*. Addison-Wesley, Reading, Mass., 1984.

[Bundy 83] A. Bundy. *The Computer Modeling of Reasoning*. Academic Press, New York, 1983.

[Bundy 85] A. Bundy, B. Silver, and D. Plummer. An analytical comparison of some rule-learning programs. *Artificial Intelligence* 27(2):137–181, November, 1985.

[Butcher 73] B. T. Butcher. *Human Intelligence*. Harper and Row, New York, 1973.

[Butter 68] C. M. Butter. *Neurophysiology: The Study of Brain and Behavior*. Brooks/Cole Publishing Co., Belmont, Calif., 1968.

[Caplan 82] D. Caplan and N. Chomsky. Language development. In D. Caplan (editor), *Biological Studies of Mental Process*. MIT Press, Cambridge, Mass., 1982.

BIBLIOGRAPHY

[Carroll 56] J. B. Carroll. *Language and Thought.* Prentice-Hall, Englewood Cliffs, N.J., 1956.

[Castleman 79] K. R. Castleman. *Digital Image Processing*. Prentice-Hall, Englewood Cliffs, N.J., 1979.

[Chang 73] C. Chang and R. Lee. *Symbolic Logic and Mechanical Theorem Proving*. Academic Press, New York, 1973.

[Changeux 85] J.- P. Changeux. *Neuronal Man*. Pantheon Books, New York, 1985.

[Cheeseman 85] Peter Cheeseman. In defense of probability. *Proc. of the Ninth International Joint Conference on Artificial Intelligence,* August 1985, Los Angeles. Calif. Morgan Kaufmann, Los Altos, Calif.

[Chomsky 75] N. Chomsky. *Reflections on Language.* Random House, New York, 1975.

[Clocksin 81] W. F. Clocksin and C. S. Mellish. *Programming in Prolog*. Springer-Verlag, Berlin, 1981.

[Cohen 82] P. R. Cohen and E. A. Feigenbaum (editors). *Handbook of Artificial Intelligence*. William Kaufmann, Los Altos, Calif., 1982. Volume 3, chapters 12 and 13.

[Cornsweet 70] T. N. Cornsweet. *Visual Perception.* Academic Press, New York, 1970.

[Courant 41] R. Courant and H. Robbins. *What is Mathematics?* Oxford University Press, London, 1941.

[Cox 46] R. Cox. Probability, frequency, and reasonable expectation. *American Journal of Physics* 14:1–13, 1946.

[Crane] H. Crane. *Beyond the Seventh Synapse*. SRI International, Menlo Park, Calif. (forthcoming).

[Crossley 72] J. N. Crossley, C. J. Ash, C. J. Brickhill, J. C. Stillwell, and N.H. Williams. *What is Mathematical Logic?* Oxford University Press, Oxford, England, 1972.

[Cunningham 72] M. Cunningham. *Intelligence: Its Organization and Development*. Academic Press, New York, 1972.

[Davis 82] R. Davis and D. B. Lenat. *Knowledge-Based Systems in Artificial Intelligence.* McGraw-Hill, New York, 1982.

[Dawkins 76] R. Dawkins. *The Selfish Gene*. Oxford University Press, New York, 1976.

[De Groot 65] A. D. De Groot. *Thought and Choice in Chess*. Basic Books Inc., New York, 1965.

[De Kleer 84] J. De Kleer. A qualitative physics based on confluences. *Artificial Intelligence* 24:7–84, 1984.

[De Kleer 84a] J. De Kleer. How circuits work. *Artificial Intelligence* 24:205–280, 1984.

[Dennett 78] D. C. Dennett. *Brainstorms: Philosophical Essays on Mind and Psychology*. Bradford Books, Montgomery, 1978.

[Dennett 84] D. C. Dennett. Cognitive wheels: the frame problem of AI. In C. Hookway (editor), *Minds, Machines, and Evolution*. Cambridge University Press, Cambridge, England, 1984.

[Dennett 84a] D. C. Dennett. *Elbow Room: The Varieties of Free Will Worth Wanting*. MIT Press, Cambridge, Mass., 1984.

[Dennett 85] D. C. Dennett. Can machines think? *In How We Know (Nobel Conference 20),* pages 121–145. Harper and Row, New York, 1985.

[Deregowski 74] J. B. Deregowski. Pictorial perception and culture. In *Image, Object, and Illusion*. W. H. Freeman, San Francisco, 1974.

[Dewdney 84] A. K. Dewdney. Computer Recreations. *Scientific American* 250:19–26, 1984.

[Dewdney 84] A. K. Dewdney. Computer Recreations: A computer trap for the busy beaver. *Scientific American* 251:19–23, 1984.

[Dewdney 85] A. K. Dewdney. Computer Recreations. *Scientific American* 253:23–24, September, 1985.

[Dretske 85] F. Dretske. Machines and the Mental. In Americal Philosophical Association Proceedings, pages 23-33. *American Philosophical Association,* Chicago, 1985.

[Dreyfus 79] H. L. Dreyfus. *What Computers Can't Do.* Harper & Row, New York, 1979, revised edition.

[Duda 73] R. O. Duda and P. E. Hart. *Pattern Recognition and Scene Analysis*. Wiley, New York, 1973.

[Duda 75] R. O. Duda and P. E. Hart. Use of the Hough transform to detect lines and curves in pictures. *Communications of the Association for Computing Machinery* 15:11–15, 1975.

[Duda 76] R. O. Duda, P. E. Hart, and N. J. Nilsson. Subjective Bayesian methods for rule-based inference systems. *Proc. AFIPS, National Comp. Conf.* 45:1075–1082, 1976.

[Duda 79] R.O. Duda, J. Gaschnig, P. E. Hart. Model design in the PROSPECTOR consultant system for mineral exploration. In D. Michie (editor), *Expert Systems in the Micro-electronics Age*. Edinburgh University, Edinburgh, 1979.

[Duda 83] R. O. Duda and Edward H. Shortliffe. Expert

System Research. *Science* 220:261–268, 1983.

[Duke-Elder 58] W. S. Duke-Elder. *System of Ophthalmology (Vol. 1, The Eye in Evolution).* C. V. Mosby Co., St. Louis, 1958.

[Dyson 79] F. Dyson. *Disturbing the Universe*. Harper and Row, New York, 1979.

[Earley 70] J. Earley. An efficient context-free parser. *Communications of the Association for Computing Machinery* 13:94–102, 1970.

[Eccles 70] J. C. Eccles. *Facing Reality*. Springer-Verlag, New York and Berlin, 1970.

[Eimas 75] P. D. Eimas. Infant perception: from sensation to perception. In L. B. Cohen and P. Salapatek (editors), *Developmental Studies of Speech Perception.* Academic Press, New York, 1975.

[Epstein 84] R. Epstein. *Bird Insight.* Nature, March, 1984. Cited in Science News, March 17, 1984, "Tracking bird-brained insight," page 171.

[Ernst 69] G. W. Ernst and A. Newell. GPS: A Case Study In Generality And Problem Solving. *Academic Press*, New York, 1969.

[Evans 68] T. G. Evans. A program for the solution of a class of geometric analogy intelligent-test questions. In M. Minsky (editor) *Semantic Information Processing*. MIT Press, Cambridge, Mass., 1968.

[Eyseneck 81] H. J. Eyseneck and L. Kamin. *The Intelligence Controversy*. Wiley, New York, 1981.

[Falmagne 75] R. J. Falmagne (editor). *Reasoning: Representation and Process in Children*. Lawrence Erlbaum Associates, Hillsdale, N.J., 1975.

[Feldman 85] J. A. Feldman. Connectionist models and parallelism in high level vision. *Computer Graphics and Image Processing* 31:178–200, 1985.

[Feller 50] B. W. Feller. *Probability Theory and Its Applications*. Wiley, New York, 1950.

[Fikes 85] R. Fikes and T. Kehler. The role of frame-based representation in reasoning. *Communications of the ACM* 28(9):904–920, Sept. 1985.

[Findler 79] N.V. Findler. *Associative Networks: The Representation and Use of Knowledge by Computers*. Academic Press, New York, 1979.

[Fingerman 69] M. Fingerman. *Animal Diversity*. Holt, Rinehart, and Winston, New York, 1969.

[Fischler 73] M. A. Fischler and R. A. Elschlager. The representation and matching of pictorial structures. *IEEE Transactions on Computers* 22:68–92, January 1973.

[Fischler 81] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24:381–395, 1981.

[Fischler and Barrett 80] Fischler, M.A. and P. Barrett, "An Iconic Transform for Sketch Completion and Shape Abstraction," Computer Graphics and Image Processing, 13:334–360, 1980.

[Flavell 63] J. H. Flavell. *The Developmental Psychology of Jean Piaget.* D. Van Nostrand, New York, 1963.

[Flavell 85] J. H. Flavell. *Cognitive Development.* Prentice-Hall Inc., Englewood Cliffs, N.J., 1985, 2nd Edition.

[Fodor 64] J. A. Fodor and J. J. Katz. *The Structure of Language.* Prentice-Hall, Englewood Cliffs, N.J., 1964.

[Fodor 74] J. A. Fodor, T. G. Bever, and M. F. Garrett. *The Psychology of Language.* McGraw-Hill, New York, 1974.

[Fodor 75] J. A. Fodor. *The Language of Thought.* Crowell, New York, 1975.

[Fodor 83] J. A. Fodor. *The Modularity of Mind.* MIT/ Bradford Books, Cambridge, Mass., 1983.

[Forgus 66] R. H. Forgus. *Perception*. McGraw-Hill, New York, 1966.

[Forgy 77] C. Forgy and J. McDermott. OPS, a domain-independent production system language. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence,* pages 933–939. Cambridge, Mass. August, 1977. William Kaufmann, Los Altos, Calif., 1977.

[Fox 82] R. Fox and C. McDaniel. The perception of biological motion by human infants. In *Science* 218:486–487, 1982.

[Frisby 80] J.P. Frisby. *Seeing—Illusions, Brain and Mind.* Oxford, Oxford, England, 1980.

[Fu 74] K. S. Fu. *Syntactic Methods in Pattern Recognition*. Academic Press, New York, 1974.

[Funt 80] B. V. Funt. Problem-solving with diagrammatic representations. *Artificial Intelligence* 13:201–230, 1980.

[Gallup 77] G. Gallup. Self-recognition in primates: a comparative approach to bidirection properties of consciousness. *American Psychologist* 32:329–338, 1977.

[Gardner 63] E. Gardner. *Fundamentals of Neurology*, W. B. Saunders Co., Philadelphia, 1963, 4th edition.

[Gardner 69] A. R. Gardner and B. T. Gardner. Teaching sign language to a chimpanzee. *Science* July 15, 1969.

[Gardner 82] H. Gardner. *Art, Mind, and Brain*. Basic Books Inc./Harper, New York, 1982. Chapter 26.

[Gardner 83] H. Gardner. *Frames of Mind: The Theory of Multiple Intelligences*. Basic Books, New York, 1983.

[Gardner 85] H. Gardner. *The Mind's New Science*. Basic Books, New York, 1985.

[Garey 79] M. R. Garey and D. S. Johnson. *Computers and Tractability*. W. H. Freeman, San Francisco, 1979.

[Garvey 81] T. D. Garvey, J. D. Lowrance, and M. A. Fischler. An inference technique for integrating knowledge from disparate sources. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, Vancouver, August, 1981. William Kaufmann, Los Altos, Calif. 1981.

[Gazzaniga 85] M. S. Gazzaniga. *The Social Brain*. Basic Books, New York, 1985.

[Gedamer 76] H.- G. Gedamer. *Philosophical Hermeneutics*. University of Calif. Press, Berkeley, Calif., 1976.

[Gentner 83] D. Gentner and A. L. Stevens (editors). *Mental Models*. Lawrence Erlbaum Associates, Hillsdale, N.J., 1983.

[Georgeff 84] M.P. Georgeff. An expert system for representing procedural knowledge. *In Joint Services Workshop on Artificial Intelligence in Maintenance, Volume 1*, Boulder, Color., 1983.

[Geschwind 79] N. Geschwind. Specializations of the brain. *Scientific American* 243:180–199, 1979.

[Gibson 50] J. J. Gibson. *The Perception of the Visual World*. Houghton-Mifflin, Boston, 1950.

[Gibson 66] J. J. Gibson. *The Senses Considered as Perceptual Systems*. Houghton Mifflin, Boston, 1966.

[Gibson 79] J. J. Gibson. *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, Mass., 1979.

[Gilchrist 79] A. L. Gilchrist. The perception of surface blacks and whites. *Scientific American* 240:88–97, 1979.

[Gloess 82] P. Y. Gloess. *Understanding LISP*. Alfred Publishing, Sherman Oaks, Calif., 1982.

[Gombrich 61] E. H. Gombrich. *Art and Illusion: A Study in the Psychology of Pictorial Presentation*. Princeton University Press, Princeton, N.J., 1961.

[Gonzalez 77] R. C. Gonzalez and P. Wintz. *Digital Image Processing*. Addison-Wesley, Reading, Mass., 1977.

[Gould 81] S. J. Gould. *The Mismeasure of Man*. W. W. Norton, New York, 1981.

[Greenberg 63] J. H. Greenberg. Some universals of grammar with particular reference to the order of meaningful elements. In J. H. Greenberg (editor), *Universals of Language*. MIT Press, Cambridge, Mass., 1963.

[Gregory 70] R.L. Gregory. *The Intelligent Eye*. McGraw-Hill, New York, 1970.

[Gregory 78] R. L. Gregory. *Eye and Brain: The Psychology of Seeing*. McGraw-Hill, New York, 1978. revised edition.

[Griffin 84] D. Griffin. *Animal Thinking*. Harvard University Press, Cambridge, Mass., 1984.

[Grosz 80] B. J. Grosz. Utterance and objective: issues in natural language communication. *AI Magazine*. 1:11–20, 1980.

[Grosz 82] B. J. Grosz, N. Haas, G. Hendrix, J. Hobbs, P. Martin, R.Moore, J. Robinson, and S. Rosenschein. *DIALOGIC: A Core Natural-Language Processing System*. Technical Report Technical Note No. 270, SRI International, Menlo Park, Calif. 94025, November 1982.

[Guilford 67] J. P. Guilford. *The Nature of Human Intelligence*. McGraw-Hill, New York, 1967.

[Gunderson 85] K. Gunderson. *Mentality and Machines*. University of Minnesota Press, Minneapolis, Minn., 1985. 2nd Edition.

[Guzman 69] A. Guzman. Decomposition of a visual scene into three-dimensional bodies. In A. Grasseli (editor), *Automatic Interpretation and Classification of Images*. Academic Press, New York, 1969.

[Hanson 78] A. R. Hanson and E. M. Riseman (eds). *Computer Vision Systems*. Academic Press, New York, 1978.

[Harmon 73] L. D. Harmon. The recognition of faces. *Scientific American* 229:71–82, 1973.

[Haugeland 85] J. Haugeland. *Artificial Intelligence: The*

Very Idea. MIT Press/Bradford, Cambridge, 1985.

[Haugeland 85a] J. Haugeland. *Mind Design*. MIT Press, Cambridge, Mass., 1985.

[Hayes 77] J. R. Hayes. Psychological differences among problem isomorphs. In D. B. Castellan, D. B. Pisoni, and G. R. Potts (editors). *Cognitive Theory*. Lawrence Erlbaum Associates, Hillsdale, N.J., 1977.

[Hayes 85] P. Hayes. The second naive physics manifesto. In: *Formal Theories of the Commonsense World*. J. Hobbes and R. Moore (editors). Ablex., Norwood, N.J., 1985.

[Hayes-Roth 83] F. Hayes-Roth, D. Waterman, and D. Lenat. *Building Expert Systems*. Addison-Wesley, Reading, Mass., 1983.

[Hegner 59] R. W. Hegner and K. A Stiles. *College Zoology*, Macmillan Co., New York, 1959, 7th edition.

[Heidegger 62] M. Heidegger. *Being and Time*. Harper and Row, New York, New York, 1962.

[Held 72] R. Held and W. Richards (editor). *Perception: Mechanisms and Models*. W.H. Freeman, San Francisco, 1972.

[Held 74] R. Held and W. Richards (editors). *Image, Object, and Illusion*. W. H. Freeman, San Francisco, 1974.

[Held 76] R. Held and W. Richards (editors). *Recent Progress in Perception*. W. H. Freeman, San Francisco, 1976.

[Hendrix 78] G. G. Hendrix. Developing a natural language interface to complex data. *ACM Transactions Database Systems* 3:105–147, 1978.

[Herbrand 30] J. Herbrand. Recherches sur la theorie de la demonstration. *Trav. Soc. Sci. Lettres Varsovie, Classe III Sci. Math. Phys.* (33), 1930.

[Herrnstein 73] R. J. Herrnstein. *IQ in the Meritocracy*. Atlantic, Little, Brown, & Co., Boston, Mass., 1973.

[Herskovits 85] A. Herskovitz. Semantics and pragmatics of locative expressions. *Cognitive Science* 9:341–378, 1985.

[Hilgard 75] E. R. Hilgard, R. C. Atkinson, and R. L. Atkinson. *Introduction to Psychology*. Harcourt Brace Jovanovich, Inc., New York, 1975, 6th edition.

[Hinchman 83] J. H. Hinchman and M.C. Morgan. Application of artificial intelligence to equipment maintenance. In *Proceedings of Joint Services Workshop on Artificial Intelligence in Maintenance. Volume 1.* Boulder, Colo., 1983.

[Hobbs 85] J. R. Hobbs. *Formal Theories of the Commonsense World*. Ablex Publishing Corp., Norwood, N.J., 1985.

[Hochberg 78] J. Hochberg. *Perception*. Prentice-Hall, Engelwood Cliffs, N.J., 1978.

[Hodges 83] A. Hodges. *Alan Turing: The Enigma*. Simon and Shuster, New York, 1983.

[Hoffman 85] D. D. Hoffman and W. Richards. Parts of recognition. *Cognition* 18:65–96, 1985.

[Hofstadter 79] D. R. Hofstadter. *Godel, Escher, and Bach*. Basic Books, New York, 1979.

[Hofstadter 83] D. R. Hofstadter and D. C. Dennett (editors). *The Mind's I*. Basic Books, New York, 1983.

[Holland 83] J. H. Holland. Escaping brittleness. In *Proceedings of the International Machine Learning Workshop*. University of Illinois, Urbana, Ill., 1983.

[Hookway 84] C. Hookway (editor). *Minds, Machines, and Evolution*. Cambridge University Press, Cambridge, England, 1984.

[Horn 74] B. K. P. Horn. Determining lightness from an image. *Computer Graphics and Image Processing* 3:277–299, 1974.

[Horn 77] B. K. P. Horn. Understanding image intensities. *Artificial Intelligence* 8:201–231, 1977.

[Horvitz 86] E. Horvitz and D. Heckerman. A framework for comparing alternative formalisms for plausible reasoning. In *AAAI-86*. Philadelphia, Penn. 1986.

[House 63] R. W. House and T. Rado. An approach to artificial intelligence. In *Proc. IEEE Winter General Meeting*, pages 111–116. IEEE, February, 1963.

[Hubel 78] D. H. Hubel. Anatomical demonstration of orientation columns in macaque monkey. *Journal of Comparative Neurology* 177:361–380, 1978.

[Hubel 79] D. H. Hubel and T. N. Wiesel. Brain mechanisms of vision. *Scientific American* 241:150–162, September, 1979.

[Hueckel 73] M. Hueckel. A local visual operator which recognizes edges and lines. *Journal of the Assoc. Comput. Machinery* 20:634, 1973.

[Hummel 83] R. A. Hummel and S. W. Zucker. On the foundations of relaxation labeling processes. *IEEE Trans. on Pattern Analysis and Machine Intelligence* PAMI-5:267–287, 1983.

[Hutchins 66] R. E. Hutchins. *Insects*. Prentice-Hall,

BIBLIOGRAPHY

Englewood Cliffs, N. J., 1966.

[Jacobs 81] G. H. Jacobs. *Comparative Color Vision.* Academic Press, New York, 1981.

[Jaynes 77] J. Jaynes. *The Origin of Consciousness in the Breakdown of the Bicameral Mind.* Houghton-Mifflin, Boston, Mass., 1977.

[Johnson-Laird 77] P. N. Johnson-Laird and P.C. Wason. A theoretical analysis of insight into a reasoning task. In P. N. Johnson-Laird and P. C. Wason (editors), *Thinking: Readings in Cognitive Science.* Cambridge University Press, Cambridge, England, 1977.

[Jolly 85] A. Jolly. A new science that sees animals as conscious beings. *Smithsonian* 15:65–75, 1985.

[Jones 74] J. R. Jones. Recursive Undecideability. *American Mathematics Monthly* 81:724–738, 1974.

[Julesz 71] B. Julesz. *Foundations of Cyclopean Perception.* University of Chicago Press, Chicago, Illinois, 1971.

[Julesz 74] B. Julesz. Cooperative phenomena in binocular depth perception. *American Scientist* 62:32–43, 1974.

[Julesz 75] B. Julesz. Experiments in the visual perception of texture. *Scientific American* 232:34–43, 1975.

[Julesz 83] B. Julesz and J. R. Bergen. Textons, the fundamental elements in preattentive vision and perception of textures. *Bell System Technical Journal* 62:1619–1645, 1983.

[Kahneman 82] D. Kahneman and A. Tversky. On the study of statistical intuitions. *Cognition* 11:123–141, 1982.

[Kanade 80] T. Kanade. Region segmentation: signal vs. semantics. *Computer Graphics and Image Processing* 13:279–297, 1980.

[Kaniza 79] G. Kaniza. *Organization in Vision.* Praeger Publishers, New York, 1979.

[Kaufman 74] L. Kaufman. *Sight and Mind.* Oxford, Oxford, England, 1974.

[Kay 73] M. Kay. The MIND system. In R. Rustin (editor), *Natural Language Processing.* Prentice-Hall, Englewood Cliffs, N.J., 1973.

[Kelly 55] G. A. Kelly. *A Theory of Personality.* W. W. Norton, New York, 1955.

[Kemeny 55] J. Kemeny. Man viewed as a machine. *Scientific American* 192:58–67, 1955.

[Kender 78] J. P. Kender. Shape from texture: a brief overview and a new aggregation transform. In *Image Understanding, Proceedings of a Workshop.* Science Applications, Washington, 1978.

[King 63] B. G. King and M. J. Showers. *Human Anatomy and Physiology.* W. B. Saunders Co., Philadelphia, 1963.

[Kline 80] M. Kline. *Mathematics, The Loss of Certainty.* Oxford University Press, New York, 1980.

[Knuth 73] D. Knuth. *The Art of Computer Programming: Fundamental Algorithms.* Addison-Wesley, Reading, Mass., 1973.

[Koestler 69] A. Koestler. *The Act of Creation.* Dell, New York, 1969.

[Kondratov 69] A. Kondratov. *Sounds and Signs.* MIR Publishers, Moscow, USSR, 1969.

[Kosslyn 80] S. M. Kosslyn. *Image and Mind.* Harvard University Press, Cambridge, Mass., 1980.

[Kosslyn 83] S. M. Kosslyn. *Ghosts in the Mind's Machine.* W. W. Norton & Co., New York, 1983.

[Kowalski 79] R. Kowalski. *Logic for Problem Solving.* North Holland, New York, 1979.

[Kuffler 76] S. W. Kuffler and J. G. Nicholls. *From Neuron to Brain.* Sinauer Associates, Inc., Sunderland, Mass., 1976.

[Kuipers 84] B. Kuipers. Commonsense reasoning about causality: deriving behavior from structure. *Artificial Intelligence* 24:169-204, 1984.

[Kripke 72] S. Kripke. Naming and necessity. In D. Davidson and G. Harmon (editors). *Semantics of Natural Language.* Reidell, Dordrecht/Boston, 1972.

[Land 59] E. H. Land. Experiments in color vision. *Scientific American* 200:84–99, 1959.

[Land 77] E. H. Land. The retinex theory of colour vision. *Scientific American* 237:108–128, 1977.

[Lane 85] H. Lane. *When the Mind Hears: A History of the Deaf.* Random House, New York, 1985.

[Larkin 81] J. Larkin. A model for learning to solve textbook physics problems. In J. R. Anderson (editor), *Skills and Their Acquisition.* Lawrence Erlbaum Associates, Hillsdale, N. J., 1981.

[Lenat 77] D. Lenat. On automated scientific theory formation: a case study using the AM program, pages 251–286. In *Machine Intelligence.* Halstead Press, New York, 1977, Volume 9.

[Lenat 84] D. B. Lenat and J. S. Brown. Why AM and Eurisko appear to work. *Artificial Intelligence* 23:269–294, 1984.

[Lenneberg 67] E. H. Lenneberg. *Biological Foundations of Language*. Wiley, New York, 1969.

[Lettvin 59] J. Y. Lettvin, R. R. Maturana, W. S. McColloch. What the frog's eye tells the frog's brain. *Proc. Inst. Radio Engineers* 47:1940–1951, 1959.

[Levine 85] M. Levine. *Vision in Man and Machine*. McGraw-Hill, New York, 1985.

[Linden 74] E. Linden. *Apes, Man, and Language*. Saturday Review Press, New York, 1974.

[Linden 86] E. Linden. *Silent Partners: The Legacy of the Ape Language Experiments*. Times Books, New York, 1986.

[Lindsay 72] P. H. Lindsay and D. A. Norman. *Human Information Processing*. Academic Press, New York, 1972.

[Lipkin 70] B. S. Lipkin and A. Rosenfeld (editors). *Picture Processing and Psychopictorics*. Academic Press, New York, 1970.

[Lowe 85] D. Lowe. *Perceptual Organization and Visual Recognition*. Kluwer, Boston, Mass., 1985.

[Luria 73] A. R. Luria. *The Working Brain: An Introduction to Neuropsychology*. Basic Books, New York, 1973.

[Lythgoe 79] J. N. Lythgoe. *The Ecology of Vision*. Clarendon Press, Oxford, England, 1979.

[MacLean 73] P. D. Maclean. *A Triune Concept of Brain and Behavior*. University of Toronto Press, Toronto, Canada, 1973.

[MacWorth 73] A. K. MacWorth. Interpreting pictures of polyhedral scenes. *Artificial Intelligence* 4:121–137, 1973.

[Malotki 83] E. Malotki. Hopi time: a linguistic analysis of the temporal concept in the Hopi language. *Trends in Linguistics, Studies and Monographs*. Number 20, Mouton, Berlin, 1983.

[Mandelbrot 82] B. B. Mandelbrot. *The Fractal Geometry of Nature*. Freeman, San Francisco, Calif., 1982.

[Mandler 83] J. M. Mandler. Representation. In J. H. Flavell and E. M. Markman (editor), *Handbook of Child Psychology: Cognitive Development*. Wiley, New York, 1983.

[Marcus 80] M. P. Marcus. *A Theory of Syntactic Recognition for Natural Language*. MIT Press, Cambridge, Mass., 1980.

[Marr 76] D. Marr. Early processing of visual information. *Philosophical Transactions of the Royal Society* 275(Series B):483–524, 1976.

[Marr 76] D. Marr and T. Poggio. Co-operative computation of stereo disparity. *Science* 194:283–287, 1976.

[Marr 78] D. Marr and H. K. Nishihara. Visual information processing: artificial intelligence and the sensorium of sight. *Technology Review* 81:1–23, 1978.

[Marr 82] D. Marr. *Vision*. W. H. Freeman, San Francisco, 1982.

[Marshall 82] J. C. Marshall. Biology of Language Acquisition. In D. Caplan (editor), *Biological Studies of Mental Processes*. MIT Press, Cambridge, Mass., 1980 (paperback edition 1982).

[Martin 78] W. N. Martin and J. K. Aggarwal. Dynamic scene analysis. *Computer Graphics and Image Processing* 7:356–374, 1978.

[McCarthy 80] J. McCarthy. Circumscription—A form of non-monotonic reasoning. *Artificial Intelligence* 13:27–39, 1980.

[McDermott 80] J. McDermott and C. Forgy. R1: An expert in the computer systems domain. In *Proceedings of the First Annual National Conf. on Artificial Intelligence*, pages 269–271, Stanford, Calif. August 1980, American Association for Artificial Intelligence, Menlo Park, Calif. 1980.

[McDermott 81] J. McDermott. R1: The Formative Years. *AI Magazine* 2:21–29, 1981.

[McKeithen 81] K. B. McKeithen. Knowledge organization and skill differences in computer professionals. *Cognitive Psychology* 13:307–325, 1981.

[McKim 72] R. McKim. *Experiences in Visual Thinking*. Brooks/Cole, Monterey, Calif., 1972.

[McNaught 75] A. B. McNaught and R. Callander. *Illustrated Physiology*. Churchill Livingstone, Edinburgh, 1975. 3rd edition.

[Miller 83] J. Miller. *States of Mind*. Pantheon, New York, 1983.

[Milne 72] L. Milne and M. Milne. *The Senses of Animals and Men*. Atheneum, New York, 1972.

[Minnaert 54] M. Minnaert. *Light and Color*. Dover, New York, 1954.

[Minsky 67] M. Minsky. *Computation: Finite and Infinite*.

Prentice Hall, Englewood Cliffs, N.J., 1967.
[Minsky 68] M. Minsky. Matter, mind and models. In M. Minsky (editor) *Semantic Information Processing*. MIT Press, Cambridge, Mass., 1968.
[Mitchell 83] T. M. Mitchell, P. E. Utzoff, and R. B. Banerji. Learning by experimentation: acquiring and refining problem-solving heuristics. In R. S. Michalski, J. G. Carbonell, T. M. Mitchell (editors), *Machine Learning, An Artificial Intelligence Approach*. Tioga Press, Palo Alto, Calif., 1983.
[Mitchell 83] T. M. Mitchell. Learning and problem solving. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, Karlsruhe, West Germany, August 1983, William Kaufmann, Los Altos, Calif., 1983.
[Moore 81] R. C. Moore. Problems in logical form. In *Proceedings of the 19th Annual Meeting of the Association for Computational Linguistics*. Stanford, Calif., June-July 1981.
[Moore 85] R. C. Moore. A formal theory of knowledge and action. In R. Hobbs and R. C. Moore (editors). *Formal Theories of the Common Sense World*. Ablex Publishing, New York, 1985.
[Mowskowitz 76] A. Mowskowitz. *The Conquest of Will—Information Processing in Human Affairs*. Addison-Wesley, Reading, Mass., 1976.
[Mueller 66] C. G. Mueller and M. Rudolph. *Light and Vision*. Time Inc., New York, 1966.

[Nagel 58] E. Nagel and J. R. Newman. *Gödel's Proof*. New York University Press, New York, 1958.
[Nagel 74] T. Nagel. What is it like to be a bat? *Philosophical Review* 83:435–451, 1974.
[Navatia 82] R. Navatia. *Machine Perception*. Prentice-Hall, Engelwood Cliffs, N.J., 1982.
[Neisser 67] U. Neisser. *Cognitive Psychology*. Appleton-Century-Croft, New York, 1967.
[Nelson 73] K. E. Nelson, G. Carskaddon, and J. Bonvillian. Syntax acquisition: impact of experimental variation in adult verbal interaction with the child. *Child Development* 44:497–504, 1973.
[von Neumann 56] cited in: *Theory of Self-Reproducing Automata*. A. W. Burks (editor). University of Illinois, Urbana, Ill., 1956.
[von Neumann 56a] J. von Neumann. Probabilistic logics and the synthesis of reliable organs from unreliable components. In *Automata Studies Number 34*. Princeton University Press, Princeton, N.J., 1956.
[Newell 72] A. Newell and H. Simon. *Human Problem Solving*. Prentice Hall, Englewood Cliffs, N.J., 1972.
[Newell 73] A. Newell. Production systems: Models of control structures. In W. G. Chase (editor) *Visual Information Processing*. Academic Press, New York, 1973.
[Newman 82] E. A. Newman and P. H. Hartline. The infrared "vision" of snakes. *Scientific American* 246:116–127, 1982.
[Nilsson 71] N. J. Nilsson. *Problem-Solving Methods in Artificial Intelligence*. McGraw-Hill, New York, 1971.
[Nilsson 80] N. J. Nilsson. *Principles of Artificial Intelligence*. Tioga, Palo Alto, Calif., 1980.
[Nilsson 83] N. J. Nilsson. Artificial intelligence prepares for 2001. *AI Magazine* 4(4), Winter 1983.
[Norman 81] D. A. Norman (editor). *Perspectives in Cognitive Science*. Ablex Publishing, Norwood, N.J., 1981.
[North 76] R.C. North. *The World That Could Be*. The Portable Stanford, Stanford Alumni Association, Stanford, Calif., 1976.

[Ohlander 78] R. Ohlander, K. Price, and D. R. Reddy. Picture segmentation using a recursive splitting method. *Computer Graphics and Image Processing* 8:313–333, 1978.
[Olds 67] J. Olds. Emotional Centers in the Brain. *Science Journal* 3:87–92, Associated Lliffe Press, London, 1967.
[Ornstein 73] R. E. Ornstein. *The Nature of Human Consciousness*. W. H. Freeman, San Francisco, 1973.

[Parzen 60] E. Parzen. *Modern Probability Theory*. Wiley, New York, 1960.
[Patterson 78] F. Patterson. Conversations with a gorilla. *National Geographic*, October 1978.
[Pavlidis 77] T. Pavlidis. *Structural Pattern Recognition*. Springer-Verlag, Berlin, 1977.
[Pavlidis 82] T. Pavlidis. *Graphics and Image Processing*. Computer Science Press, Rockville, 1982.
[Pearl 84] J. Pearl. *Heuristics: Intelligent Strategies for Computer Problem Solving*. Addison-Wesley, Reading, Mass., 1984.

[Penfield 78] W. Penfield. *The Mystery of the Mind.* Princeton University Press, Princeton, N.J., 1978.

[Pentland 83] A. P. Pentland and M. A. Fischler. A more rational view of logic. *AI Magazine* 4(4), Winter 1983.

[Pettigrew 72] J. D. Pettigrew. The neurophysiology of binocular vision. *Scientific American* 227:84–95, 1972.

[Piaget 52] J. Piaget. *The Origins of Intelligence in Children.* International Libraries Press, New York, 1952.

[Pirenne 67] M. H. Pirenne. *Vision and the Eye.* Barnes and Noble, New York, 1967.

[Polya 54] G. Polya. *Patterns of Plausible Inference.* Princeton University Press, Princeton, N.J., 1954.

[Popper 77] K. R. Popper and J. C. Eccles. *The Self and the Brain.* Springer-Verlag, London, 1977.

[Pratt 76] W. K. Pratt. *Digital Image Processing.* Academic Press, New York, 1976.

[Pribram 71] K.H. Pribram. *Languages of the Brain.* Prentice-Hall, Englewood Cliffs, New Jersey, 1971.

[Putnam 75] H. Putnam. The meaning of meaning. In H. Putnam (editor) *Mind, Language, and Reality: Philosophical papers*, Vol. 2, Cambridge University Press, Cambridge, England, 1975.

[Pylyshyn 78] Z. W. Pylyshyn. Imagery and artificial intelligence. In Ned Block (editor), *Readings in Philosophy of Psychology.* Harvard University Press, Cambridge, Mass., 1981.

[Pylyshyn 73] Z. W. Pylyshyn. What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin* 80:1–24, 1973.

[Pylyshyn 80] Z. W. Pylyshyn. The causal power of machines. *The Behavioral and Brain Sciences* 3:442–444, 1980.

[Rado 62] T. Rado. On non-computable functions. *Bell Systems Technical Journal* 41:877–884, 1962.

[Raphael 76] B. Raphael. *The Thinking Computer.* W. H. Freeman, San Francisco, 1976, chapter 6.

[Ratliff 72] F. Ratliff. Contour and contrast. *Scientific American* 226:90–101, 1972.

[Restak 79] R. Restak. *The Brain, The Last Frontier.* Warner Books, New York, 1979.

[Restak 84] R. M. Restak. *The Brain.* Bantam Books, New York, 1984.

[Reynolds 78] A. G. Reynolds and P. W. Flagg. *Cognitive Psychology.* Winthrop, Cambridge, Mass, 1978.

[Richardson 72] K. Richardson, D. Spear, and M. Richards (editors). *Race and Intelligence.* Penguin Books, Baltimore, 1972.

[Richie 84] G. D. Richie and F. K. Hanna. AM: A case study in AI methodology. *Artificial Intelligence* 23:249–268, 1984.

[Rips 77] L. J. Rips. Supposition and the analysis of conditional sentences. In M. A. Just and P. A. Carpenter (editors) *Cognitive Processes in Comprehension.* Lawrence Erlbaum Associates, Hillsdale, N.J., 1977.

[Robinson 65] J. A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM* 12:23–41, 1965.

[Robinson 73] D. N. Robinson. *The Enlightened Machine.* Dickeson Publishing Co., Belmont, Calif., 1973.

[Robinson 82] J. J. Robinson. DIAGRAM: A grammar for dialogues. *Communications of the ACM* 25:27–47, 1982.

[Rock 85] I. Rock. *The Logic of Perception.* MIT Press, Cambridge, Mass., 1985.

[Romer 55] A. S. Romer. *The Vertebrate Body.* W. B. Saunders, Philadelphia, 1955.

[Rosch 77] E. Rosch. Classification of real world objects: origins and representations in cognition. In P. N. Johnson-Laird and P. C. Wason (editors), *Thinking: Readings in Cognitive Science.* Cambridge University Press, Cambridge, England, 1977.

[Rosenfeld 76] A. Rosenfeld and A. Kak. *Digital Picture Processing.* Academic Press, New York, 1976.

[Rosenschein 83] S. Rosenschein. Natural language processing: crucible for computational theories of cognition. In *International Joint Conference on Artificial Intelligence*, pages 1180–1186. Karlsruhe, West Germany, William Kaufmann, Los Altos, Cal., August, 1983.

[Ross 80] G. S. Ross. Categorization in 1- to 2-year-olds. *Developmental Psychology* 16:391–396, 1980.

[Sacks 86] O. Sacks. Mysteries of the deaf. *New York Review of Books*, March 27, 1986.

[Sagan 78] C. Sagan. *The Dragons of Eden.* Ballantine Books, New York, 1978.

[Salapatek 77] P. Salapatek. Stimulus determinants of attention in infants. In B. B. Wolman (editor), *International Encyclopedia of Psychiatry, Psychology, Psychoanalysis, and Neurology*. Van Nostrand Reinhold, New York, 1977.

[Samuel 67] A. L. Samuel. Some studies in machine learning using the game of checkers II: recent progress. *IBM Journal of Research and Development* 11:601–617, 1967.

[Schank 81] R. Schank and C. R. Riesback. *Inside Computer Understanding*. Lawrence Erlbaum Asociates, Hillsdale, N.J., 1981.

[Scientific 72] Scientific American. *Perception: Mechanisms and Models*. W. H. Freeman, San Francisco, Calif., 1972. (Readings from the *Scientific American*; paperback.)

[Scientific 72a] Scientific American. *Communication*. W. H. Freeman, San Francisco, 1972. (Readings from the *Scientific American*; paperback.)

[Scientific 74] Scientific American. *Object and Illusions*. W. H. Freeman, San Francisco, 1974. (Readings from the *Scientific American*; paperback.)

[Scientific 76] Scientific American. *Progress in Perception*. W. H. Freeman, San Franscisco, 1976. (Readings from the *Scientific American*; paperback.)

[Scientific 79] Scientific American. *The Brain*. W. H. Freeman and Co., San Francisco, CA., 1979. (Readings from *Scientific American*; paperback.)

[Scribner 77] S. Scribner. Modes of thinking and ways of speaking: culture and logic reconsidered. In P. N. Johnson-Laird and P. C. Wason (editors), *Thinking: Readings in Cognitive Science*. Cambridge University Press, Cambridge, England, 1977.

[Searle 69] J. R. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, London, 1969.

[Searle 83] J. R. Searle. Minds, brains, and programs. In D. R. Hofstadter and D. C. Dennett (editors) *The Mind's I*. Basic Books, New York, 1983.

[Searle 84] J. R. Searle. *Minds, Brains and Action*. Harvard University Press, Cambridge, Mass., 1984.

[Sebeok 80] T. A. Sebeok and D. J. Umiker-Seboek (editors). *Speaking of Apes: A Critical Anthology of Two-Way Communication with Man*. Plenum, New York, 1980.

[Sejnowski 86] T. J. Sejnowski and C. R. Rosenberg. *NETtalk: A Parallel Network That Learns to Read Aloud*. Tech. Report JHU/EECS-86/01, The Johns Hopkins University Electrical Engineering and Computer Science Dept., Baltimore, 1986.

[Shafer 76] G. A. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, N.J., 1976.

[Shafer 85] G. A. Shafer and A. Tversky. Weighing evidence: the design and comparison of probability thought experiments. *Cognitive Science* 9:309–339, 1985.

[Shanon 85] B. Shanon. The role of representation in cognition. In J. Bishop, J. Lockhead, D. N. Perkins (editor), *Thinking*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1985.

[Shepard 71] R. N. Shepard and J. Metzler. Mental rotations of three-dimensional objects. *Science* 171:701–703, 1971.

[Shepard 82] R. N. Shepard and L. A. Cooper. *Mental Images and Their Rotations*. MIT Press, Cambridge, Mass., 1982.

[Shifrin 77] R. M. Shifrin and W. Schneider. Controlled and automatic human information processing II: perceptual learning, automatic attending, and a general theory. *Psychological Review* 84:127–190, 1977.

[Shirai 75] Y. Shirai. Generating schematic descriptions from drawings of scenes with shadows. In P. H. Winston (editor), *Psychology of Computer Vision*. McGraw-Hill, New York, 1975.

[Shortliffe 76] E. H. Shortliffe. *Computer-based Medical Consultations: MYCIN*. Elsevier Publishing Co., New York, 1976.

[Simon 81] H. A. Simon. *Cognitive Science: The Newest Science. Perspectives on Cognitive Science*. Ablex Publishing, Norwood, N.J., 1981.

[Skinner 57] B. F. Skinner. *Verbal Behavior*. Appleton-Century-Crofts, New York, 1957.

[Sloman 85] A. Sloman. What enables a machine to understand? In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, August 1985, Morgan Kaufmann, Los Altos, Calif., 1985.

[Smith 47] S. M. Smith, H. O. Brown, J. E. P. Toman, and L. S. Goodman. The lack of cerebral effects of d-tubercurarine. *Anesthesiology* 8:1–14, 1947.

[Smith 72] C. U. M. Smith. *The Brain, Towards an Understanding*. Capricorn Books, New York, 1972.

[Smith 85] C. G. Smith. *Ancestral Voices*. Prentice-Hall, Englewood Cliffs, N.J., 1985.

[Smith 83] R. G. Smith, and J. D. Baker. The dipmeter advisor system: a case study in commercial expert system development. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, pages 122–129. Karlsruhe, West Germany, August, 1983. William Kaufmann, Los Altos, Calif.

[Soloway 80] E. M. Soloway. *From Problems to Programs Via Plans*. Technical Report COINS Technical Report 80-19, University of Mass., Amherst, Mass., 1980.

[Spelke 82] E. S. Spelke. Perceptual knowledge of objects in infancy. In J. Mehler, E. C. T. Walker, and M. Garrett (editor), *Perspectives in Mental Representation*. Lawrence Erlbaum Associates, Hillsdale, N.J., 1982.

[Spence 83] A. P. Spence and E. B. Mason. *Human Anatomy and Physiology*. Benjamin Cummings Publishing Co., Menlo Park, Calif., 1983, 2nd edition.

[Springer 85] S. P. Springer and G. Deutsch. *Left Brain, Right Brain*. W. H. Freeman, San Francisco, Calif., 1985, revised edition.

[Sternberg 82] R. J. Sternberg (ed). *The Handbook of Human Intelligence*. Cambridge University Press, New York, 1982.

[Svedja 79] M. Svedja and D. Schmid. The role of self-produced locomotion on the onset of heights on the visual cliff. *Paper presented at the meeting of the Society for Research in Child Development*, San Francisco, March, 1979.

[Szlovits 82] P. Szlovits (editor). *Artificial Intelligence in Medicine*. Westview, Boulder, Colo., 1982.

[Terrace 81] H. S. Terrace. *Nim: A Chimpanzee Who Learned Sign Language*. Knopf, New York, 1980.

[Thompson 84] H. Thompson and G. Ritchie. Implementing Natural Language. In T. O'Shea and M. Eisenstadt (editors) *Artificial Intelligence: Tools, Techniques, and Applications*. Harper and Row, New York, 1984.

[Thorpe 65] W. H. Thorpe and M. E. W. North. Origin and significance of the power of vocal imitation: with special reference to the antiphonal singing of birds. *Nature* 208:219–222, 1965.

[Tou 74] J. T. Tou and R. C. Gonzalez. *Pattern Recognition Principles*. Addison-Wesley, Reading, Mass., 1974.

[Treisman 85] Treisman, A., "Preattentive Processing in Vision," Computer Vision, Graphics, and Image Processing, 31(2):156–177, August 1985.

[Turing 50] A. A. Turing. Computing machinery and intelligence. *Mind* 59:433–460, 1950.

[Turner 85] R. Turner. *Logics for Artificial Intelligence*. Halsted/Wiley, New York, 1985.

[Tversky 74] A. Tversky and D. Kahneman. Judgment under uncertainty: heuristics and biases. *Science* 185:1124–1131, 1974.

[Ullman 79] S. Ullman. *The Interpretation of Visual Motion*. M.I.T. Press, Cambridge, Mass., 1979.

[Walls 42] G. Walls. *The Vertebrate Eye and Its Adaptive Radiation*. Cranbrook Institute of Science, Bloomfield Hills, Michigan, 1942. (also Hafner, New York 1942).

[Waltz 75] D. I. Waltz. Generating schematic descriptions from drawings of scenes with shadows. In P. H. Winston (editor), *The Psychology of Computer Vision*. McGraw-Hill, New York, 1975.

[Wang 60] H. Wang. Toward mechanical mathematics. *IBM Journal of Research and Development* 4:2–22, 1960.

[Wasserman 73] A. O. Wasserman. *Biology*. Appelton-Century-Crofts, New York, 1973.

[Watson 30] J. Watson. *Behaviorism*. W. W. Norton, New York, 1930.

[Weinreb 84] E. L. Weinreb. *Anatomy and Physiology*. Addison-Wesley, Reading, Mass., 1984.

[Weiss 79] S.M. Weiss and C.A. Kulikowski. EXPERT: A system for developing consultation models. In *Proceedings of the 6th International Joint Conference on Artificial Intelligence*, pages 942–947, Tokyo. William Kaufmann, Los Altos, Calif., 1979.

[Weissman 67] C. Weissman. *LISP 1.5 Primer*. Dickenson, Belmont, Calif.,1967.

[Weizenbaum 66] J. Weizenbaum. ELIZA, a computer program for the study of natural language communication between man and machine. *Communica-*

tions of the Association for Computing Machinery 9:36–45, 1966.

[Weizenbaum 76] J. Weizenbaum. *Computer Power and Human Reason*. W. H. Freeman and Co., San Francisco, 1976.

[Wells 68] M. Wells. *Lower Animals*. McGraw Hill, New York, 1968.

[Wertheimer 61] M. Wertheimer. *Productive Thinking*. Tavistock, London, 1961. (Original edition 1945, Harper & Bros., New York)

[Whorf 56] B. L. Whorf In J. B. Carroll (editor). *Language, Thought, and Reality: Selected Writings of Benjamin Lee Whorf*. Wiley, New York, 1956.

[Wigglesworth 72] V. B. Wigglesworth. *Principles of Insect Physiology*. Chapman and Hall, London, 1972. 7th edition.

[Wilensky 84] R. Wilensky. *LISPcraft*. W. W. Norton, New York, 1984.

[Winograd 70] T. Winograd. *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language*. PhD thesis, MIT, Cambridge, Mass., August, 1970.

[Winograd 74] T. Winograd. *Five Lectures on Artificial Intelligence*. Report STAN CS 74-459, Computer Science Department, Stanford University, September 1974.

[Winograd 83] T. Winograd. *Language as a Cognitive Process: Syntax*. Addison-Wesley, Reading, Mass., 1983.

[Winograd 86] T. Winograd and F. Flores. *Understanding Computers and Cognition*. Ablex Publishing Co., Norwood, N.J., 1986.

[Winston 75] P. H. Winston. Learning Structural Descriptions from Examples In P. H. Winston (editor). *The Psychology of Computer Vision*. McGraw-Hill, New York, 1975.

[Winston 84] P. H. Winston. *Artificial Intelligence*. Addison-Wesley, New York, 1984. 2nd edition.

[Winston 85] P. H. Winston and B. K. Horn. *LISP*. Addison-Wesley, New York, 1985. 2nd edition.

[Wittgenstein 68] L. Wittgenstein. *Philosophical Investigations*. Macmillan, New York, 1968.

[Woods 77] W. A. Woods. Lunar rocks in natural English: explanations in natural language question answering. In *Linguistic Structures Processing*. A. Zampoli (editor). North-Holland, Amsterdam, 1977.

[Wos 84] L. Wos, R. Overbeek, E. Lusk, and J. Boyle. *Automated Reasoning: Introduction and Applications*. Prentice Hall, Englewood Cliffs, N.J., 1984.

[Wright 69] D. Wright. *Deafness*. Stein and Day, New York, New York, 1969.

[Yonas 81] A. Yonas. Infant's response to optical information for collision. In R. N. Aslin, J. R. Alberts, and M. R. Peterson (editor), *Development of Perception: Psychobiological Perspectives*. Academic Press, New York, 1981.

[Young 68] J. Z. Young. *Doubt and Certainty in Science*. Oxford University Press, New York, 1968.

[Young 76] R. M. Young. *Seriation by Children: An Artificial Intelligence Analyis of a Piagetian Task*. Birkhauser, Basel, Switzerland, 1976.

[Zahn 71] C. T. Zahn. Graph-theoretical methods for detecting and describing gestalt clusters. *IEEE Transactions on Computers* C-20:68–86, 1971.

[Zucker 76] S. Zucker. Region growing: childhood and adolescence. *Computer Graphics and Image Processing* 5:382–399, 1976.

[Zucker 77] S. W. Zucker, R. A. Hummel, and A. Rosenfeld. An application of relaxation labeling to line and curve enhancement. *IEEE Transactions on Computers* C-26:394–403, 1977.

Index

INDEX