COGNITIVE SCIENCE 12, 467-483 (1988)

Book Reviews

Appelt, D.E. *Planning English Sentences*. Cambridge, England: Cambridge University Press, 1985. 171 pages.

Reviewed by *Allan Ramsay, School of Cognitive Sciences, University of Sussex, Falmer BN1 9QN, England.*

Appelt, in the work reported in this book, has done something which far too few people in AI have done. He has taken a number of existing theories and program descriptions and synthesized them into a single system (acronymed KAMP for "knowledge and modalities planner"). This is in stark contrast to the majority of work in AI, which pays lip service to the principle of basing one's work on the reports of what has gone before, but which then starts from scratch with yet another new notation, new set of principles, and new primitive notions. Appelt, by trying to build his system out of existing components, provides us with an unusual opportunity of assessing the viability of these theories as contributions to an overall explanation of language processing. Indeed, because most of them are not his own invention—and therefore he does not have the emotional investment in them that inventors of AI systems usually have—his own criticisms of them are more penetrating and insightful than is often the case in the perfunctory discussions included at the end of doctoral dissertations.

So what is the task that he wants KAMP to perform, and what are the theories he is using as a basis? The essence of his task is captured by the title of the book. He wants a system which treats the construction of English surface strings as a rational activity just like any other goal-directed activity, and which therefore accounts for this construction using planning mechanisms very similar to the ones it uses for the rest of its activities. There are, as we will see, problems associated with assessing whether the system satisfies his aims, but the task is intuitively clear. And the theories? The main levels of analysis which he recognizes are reasoning about knowledge and belief, general planning, speech act theory, grammatical constraints, and to some extent discourse theory. There is an awful lot to be covered here, and it would have been sheer madness to try to invent his own theories for each level. It would have been equally impossible to perform thorough experiments using each of the competing available theories for them—there is quite enough work in gluing together *one* theory from each level, without asking for alternatives. If anyone were to criticize Appelt's choice of theory at some level, he would be quite entitled to invite them to do their own reimplementation and then perform the relevant comparisons. The choices he did make seem to have been dictated by a variety of considerations, ranging from principled consideration of alternatives to acceptance of one because he had an implementation readily available. Whatever the reasons for the decisions, the actual theories chosen are certainly not way out of line: Moore's (1980, 1984) theory of knowledge and action, Sacerdoti's (1977) nonlinear planner, the extension of Searle's (1969) theory of speech acts by Cohen and colleagues, Kay's (1985) functional unification grammar, and the work by Sidner (1983) and Grosz, Joshi, and Weinstein (1983) on discourse phenomena. If a system based on a combination of these theories and programs has gaping flaws, so will any theory based on the obvious competitors.

The book in which this system is presented is a very slim volume, clearly a cut-down version of Appelt's thesis. This is unfortunate, since quite a common reaction among people who have read it is "But how much of it is implemented, and how is it all glued together?". We cannot get much of an answer to this. The revision of the thesis to turn it into an accessible, readable volume has chopped out so much that there is very little practical detail, and very little discussion of what the system will and won't do. There are two ways of approaching this question. We can look for internal evidence, possibly arguing that the claim in this book for KAMP is all it can do, on the ground that if it did more we would have been told so. Following this approach, it would be possible to conclude that the system's range is really quite limited, since only one example is followed through in any detail. Arguing like this, though, is very dangerous, since the one example which is talked through in detail (the generation of the sentence *Remove the pump with the wrench in the tool box*) requires 21 pages of description. To demand another example might simply have taken another 20 or so pages, without adding greatly to our understanding.

But if we cannot assess KAMP's coverage from what is actually present in the book, what can we do? The only available option is to consider the theories on which it is based and see if there is anything we know about which they cannot cover. Three things occurred to this reader, namely the question of mutual knowledge, problems about quantification in modal representations of knowledge, and problems of constructing appropriate referring expressions. As we remarked above, if these are flaws in the underlying theories then they are liable to be present in any competitor to Appelt's analysis, since these theories are among the best-developed and most sophisticated that we have. Consideration of these problems, then, should not be taken as an attack on Appelt's work, since it does seem to be an implementation of state-of-the-art theory. Nonetheless, to see what should be done next, we have to get a feel for the power of what has been done so far, and there does not seem to be any easy way to get that from the present volume. We will consider the three problems identified above one at a time, to see if there are any *a priori* limits on what can be achieved by the techniques employed in KAMP.

Mutual knowledge: It is well known that correct interpretation by a hearer of an utterance can require considerable amounts of knowledge about the state of the speaker's knowledge, and in particular of the speaker's knowledge of the hearer's knowledge, and the speaker's knowledge of the hearer's knowledge of the speaker's knowledge, and so on *ad nauseam* (we will use the terms "speaker" and "hearer" to refer respectively to the person or system that generated an utterance or text and the person or system that it was addressed to, regardless of the fact that virtually all AI systems, including Appelt's, deal with text and not speech). We see the necessity for such knowledge cropping up at virtually every level of linguistic processing as soon as we start thinking about language as a social process involving a speaker and a hearer. To refer to some tool as a "cotterless chainset extractor," I have to believe that my hearer knows what these words mean. To refer to it as "the cotterless chainset extractor," I have to have some picture of what they know about the current environment, so that I know there is one such tool which they will recognize as being particularly salient to me in the current context. For me to use a sentence like "The cotterless chainset extractor is in the other room" as a request to them to go and get it for me, I have to be able to reason that they will understand that I am telling them this fact in order that they will be able to help me by getting it for me.

What is less clear is what form this knowledge takes, or how we acquire and use it. The first problem is that there is no known limit on how deep the nesting of knowledge about the other person's knowledge needs to go. It is not too difficult to construct cases in which seven or eight layers are necessary, and although it tends to get rather intricate beyond that point there seems no principled reason to suppose that there should be a finite bound. There are not, however, any overwhelming arguments to suggest that fully mutual knowledge is ever required, and there are hints from computer science that it is not, in fact, achievable. Garrod and Anderson (1987) suggest that much of the "repair"-oriented nature of normal spoken discourse arises because although we would need such knowledge in order to come up with reliable plans incorporating speech acts, what we actually have is tentative beliefs about each other. They support their argument by analysis of dialogues between humans communicating in natural language during the course of cooperative problem solving where neither participant initially knows anything about what the other knows (the subjects play a computer-based maze exploration game, in which one's own progress can depend on persuading the other participant to move to specific locations, but where neither participant initially knows where the other is in the maze). Many of the utterances in these dialogues appear to be aimed at establishing some minimal common ground on which to base the sort of reasoning we discussed above. The notable thing about this process of generating a common ground is that it never seems to acquire the status of knowledge, tending to remain rather tentative. This seems fair enough, since the acts which establish it are themselves based on unreliable inferences from the previous common ground, which in its turn was never reliably established as common.

Appelt attempts to deal with this issue by introducing the notion of the "kernel" of two agents' knowledge. The kernel is taken to be a collection of propositions for which either agent can infer any of the required nested propositions; that is, if P is in the kernel of A and B's knowledge then either A or B may assume any formula such as knows(A, knows(B, knows(B, knows(A, P)))).

This specification of the kernel seems initially to provide what is required for reasoning appropriately about the effects and preconditions of speech acts. If, for example, a system knew that P was in the kernel for it and its hearer, then it should be able to work out that it is inappropriate for it to attempt to inform its hearer that P is the case, since you cannot properly inform someone of something you know they already know. It is this property of informing that enables you to say "The door's open" to someone who has just come in and left the door open as a way of asking them to shut it—since they know that you both already know that the door is open they can infer that you were not genuinely attempting to inform them, and hence are forced to try to work out what you really were doing. We would therefore like to see KAMP using the kernel in this way—at the very least for working out whether the basic act of informing is going to work properly, and perhaps for planning indirect acts as well.

If we look at the specification that Appelt provides for the action of informing somebody about some fact, we note several things. The given axiomatization is as follows:

Preconditions:
    ∀A,B,P,w1,w2 R(:Do(A, :Informif(B,P), w1, w2) →
        V(w1, :Location(A)) = V(w1, :Location(B)) &
        T(w1, Intend(@A, Know(@B, @P)) V Intend(@A, Know (@B,
            ¬@P))) &
        T(w1, Know(@A, @P) V Know(@A, ¬@P))

Effects:
    ∀A,B,P,w1,w2 R(:Do(A, :Informif(B,P)), w1, w2) →
        ∀w3 K(Kernel(A,B), w2, w3) →
        ∃w4 K(Kernel(A,B), w1, w4) & R(:Do(A, :Informif(B,P)), w4, w3))

Much of the baroque nature of the formalism here is due to Appelt's decision to follow Moore in dealing with modal logic by reasoning using a first-order axiomatization of the possible worlds semantics. This treatment was probably the best available at the time Appelt was working on his system, though it may be that recent work by Wallen (1987) would in fact have given a more effective theorem prover and a more readable notation. Be that as it may, we will accept Appelt's decision in the context in which he was working, ignore the fine details of the notation, and try to give a flavor of what these axioms mean and what sort of reasoning they do and do not support.

The precondition axiom is intended to capture three things. It says (very roughly) that for w2 to be a world which could be the result of A informing B of the truth or otherwise of P in a world w1, then three propositions must hold. A and B must have been in the same place in w1, A must have either wanted B to know that P was true or wanted him to know that it was false, and A must have actually known whether P was true. The axiom as given in the book has

T(w1, Know(@B, @P) V Know(@B, ¬@P))

as the last line of the preconditions, in other words that for A to inform B about the truth or otherwise of P B must already know whether P is true or not. I assume that this is a typo.

The axiom for the effect of the action says (equally roughly) that if w2 is the sort of world that would result from A informing B about P in world w1, then in w2 both parties mutually know that A has indeed done this action.

The appealing thing about this pair of axioms is that the effect axiom makes no mention of anything other than that they both subsequently know that A did indeed inform B about P. This in itself gives them immediate access to all sorts of interesting facts, since for A to have performed this action then the precondition axiom must have been satisfied. Since it is presumed that all speakers have mutual access to all the axioms about speech acts and about knowledge, they both have immediate access to the truth value of P (in this tradition you cannot know things that are not true), and they also mutually know this. The attraction of this is that we seem to get a lot of work out of what looks like a fairly minimal set of axioms. The problem is that careful reflection shows firstly that the axioms are not in fact adequate for a proper characterization of informing, and secondly that they do not seem to support quite the range of inferences that might be required.

The inadequacies show up in various ways. Firstly, since the precondition axiom makes no mention of the requirement that A should know that B does not already know whether P is true, we lack the information we need for distinguishing between genuine attempts at informing and the use of declarative statements for other goals, as discussed above. Subsequent axioms for the act of describing an object do refer to mutual knowledge in their preconditions, but the axioms for informing given in the text do not, thus missing some of the point of having any analysis of mutual knowledge in the first place. Furthermore, in the absence of a more complete characterization of the preconditions it is hard to see how B could confirm that what

A had just been doing was an informing. Without this one cannot legitimately infer that the effects of an act of informing now hold (and A therefore cannot legitimately infer that they do either), so the route by which the fact that A informed B of P (or not P) gets into their kernel becomes highly unreliable.

The problem about getting these axioms to support the required range of inferences concerns the disjunctions that appear inside the precondition axiom. The first area of difficulty involves the part of the axiom that says that A cannot inform B about P unless he either knows P is true or he knows it is false. After the act of informing has been successfully carried out and accepted by both participants as a success, their kernel contains that fact that A either knew it was true or that it was false, and that it was indeed either true or false. There does not seem to be enough machinery around, however, for B to work out which was the case—that it was true and that A knew it was true, or that it was false and A knew it was false. An appeal to the fact that A knew which was the case will not do. We need some extra mechanism (which may have been present in Appelt's program, but is not clearly visible in the current axiomatization) which will enable B to make this extra move. Otherwise he or she has simply discovered that A has the required information, not what it is. The second difficulty concerns the relationship between the fact that A knew the status of P and what he or she intended to convey to B. It is far from clear that any mechanism which enabled B to work out, for instance, that what A knew was that P was true, and to infer for himself that P was indeed true, could reasonably be expected to work out from the current set of axioms that A wanted him or her to know it was true (rather than wanting him or her to know it was false). Again simply appealing to the fact that things that are known have to be true will not do. Surely part of the aim of any axiomatization like this should be to help us work out whether we are being told the truth or lied to, and to make our inferences accordingly.

These apparently nitpicking criticisms of a single pair of axioms should not be taken as an attack on the adequacy of this particular attempt to construct a set of axioms for dealing with speech acts. Appelt seems to have done as good a job as anyone else would have been likely to do. But nor should they be taken as things which could have been patched by simply tagging on a few more clauses. The extremely tentative nature of naturally occurring dialogues, together with the fact that even an apparently very carefully constructed set of axioms seems very frail when subjected to stress, should make us wary. In particular, if we want to use the axiomatizations both to help us work out what speech act was intended and to help us build the kernel set of propositions which constitutes our mutual knowledge, we will have to be even more careful than we had previously realized. In the case of informing, for instance, just adding the precondition that A should not know that B already knows whether P is true will have radical effects on A and B's supposed mutual knowledge. Suppose, for instance, that B has somehow got the impression that A does know whether he or she knows P. A may then think that he or she informed B, and may draw all the mutual knowledge inferences that this would legitimize; but since B thinks something else happened, these inferences will not be valid.

The second set of problems concerns the operation of quantifying in. It is well known that the treatment of existential quantifiers in any epistemic logic is problematical. An example due to Quine concerns the ambiguity of the sentence *Ralph knows that someone in this room is a spy*. Taking a very naive formalism, we might translate this either as $\equiv \exists x(\mathrm{here}(x)\ \&\ \mathrm{knows}(R, \mathrm{spy}(x)))$ or as $\mathrm{knows}(R, \exists x(\mathrm{spy}(x)\ \&\ \mathrm{here}(x)))$. It is clear that the intended meanings of these formalizations are different. In the first Ralph knows who the spy is, for instance because he saw them purloining documents, whereas in the second he may simply know that 10 minutes ago the documents were in his desk, that they are not there now, and that no one has left the room in the last 10 minutes. This distinction must be maintained for any representation which is to be used for planning and interpreting general linguistic acts. In particular, it is hard to see how Ralph could plan an utterance such as *I know one of you has stolen the documents. Come on now, who was it?* if $\exists x(\mathrm{here}(x)\ \&\ \mathrm{knows}(R, \mathrm{spy}(x)))$ and $\mathrm{knows}(R, \exists x(\mathrm{spy}(x)\ \&\ \mathrm{here}(x)))$ are equivalent. If they are, then in order to produce the first sentence of the given utterance the speaker must already know the answer to the question.

In the theory which Appelt has borrowed for performing reasoning about knowledge (Moore's [1984] reified modal logic) the only way to preserve any distinction between these two is to suppose that everything you might want to talk about has a name, which is known to everyone and which alone suffices as an answer to questions about identity. In particular, Appelt borrows the following axiom for dealing with existential quantification:

$$\forall w(T(w, \exists xP) \equiv \exists x T(w, P[@x/x]))$$

To see what this means we have to be a bit less cavalier about our treatment of the notation. Without going too far into the fine details, though, we can interpret it as saying that $\exists xP$ is true in some world if the formula that you would get by substituting some *rigid designator* (r.d.) for x in P is true in that world. The use of r.d.'s is critical here. An r.d. is a term whose denotation is the same in all possible worlds. @ is taken to be a function from the domain of discourse to a set of r.d.'s. R.d.'s are rather like proper names, or like proper names ought to be. Very few proper names, however, work correctly as r.d.'s, because there is usually more than one person or thing with the same name. *Socrates* and *The University of Sussex* probably fit the criteria for r.d.'s, but *Allan Ramsay* and *Paris* certainly don't—you need context to know whether we are talking about Allan Ramsay the AI lecturer, or Allan Ramsay the painter, or any of the thousands of other Allan Ramsay's there are and have been; you even need context to know whether we are talking about Paris, France, or Paris, TX, or Paris who fell in love with Helen of Troy. Furthermore, very few things have well-behaved proper names, whereas the analysis of the semantics of modal logic in terms of r.d.'s requires everything to have one.

Once we have introduced r.d.'s, we can use them to treat statements such as *There is someone who Ralph knows to be a spy.* Translating this into Appelt's notation we would get something like:

$$T(W0, \exists x\ Know(R, spy(x))) \rightarrow [\exists y \forall w1\ K(R, W0, w1) \rightarrow T(w1, spy(@y))]$$

In other words, if in the real world there is someone who Ralph knows to be a spy, then there is an individual who is a spy in every world compatible with what Ralph knows; and furthermore there is a term which denotes that individual in all those worlds.

The first part of this is uncontentious. If there is indeed someone who Ralph knows to be a spy, then if you changed the truth values of every sentence about which Ralph had no knowledge then that same person would presumably still be a spy. The problem arises from the need to asume that there is an r.d. whose value is this person. Firstly, this seems to imply that we must have an r.d. for any object that we can have knowledge about or talk about—that @ is a total function on the universe of the set of possible worlds. This just does not map onto our intuitions about our own knowledge and language, where we choose fleetingly appropriate nonrigid designators for referring to people and things exactly because we do not have names for everything. Secondly, the introduction of r.d.'s to deal with existential quantification leads insidiously to the assumption that WH-questions are always requests to be provided with an r.d. This is just plain wrong, as Appelt implicitly concedes when talking about the axioms that characterize the action of describing something.

To see that it is wrong, consider the following examples. (1) **Q:** *Who is Mrs. Smith?* **A:** *She's the mayor of Brighton.* The questioner already has something which is as near to being an r.d. as we ever get in a natural language, that is, a proper name, yet still feels that he does not have the information he needs. What he gets as an answer is a nonrigid designator. (2) *You remember we were talking about Mrs. Smith yesterday. Well that's her over there.* Again we have a proper name being clarified by the addition of extra information, this time about how to identify the person designated by the name.

Barwise and Perry (1983) argue that referring expressions can carry many different kinds of information, depending on the states of the knowledge of the speaker and hearer. Groenendijk and Stokhof (1987) argue at length about the kind of entities that constitute answers to WH-questions. It seems clear from this work that attempts to deal with the required knowledge simply in terms of r.d.'s are just not going to be rich enough to capture the required distinctions. Unfortunately, very few notations which are sufficiently rich are also computationally tractable, since it seems we are going to have to deal with intensions as well as intentions. Turner (1987) has suggested a theory of properties which may be amenable to first-order treatment, but as yet there seems to be no implemented system for dealing with this theory. Appelt's own theory of referring expressions, to which we turn next, seems to reflect an awareness that the treatment via r.d.'s is going to need modification before any utterance can actually be constructed. It seems a pity that this awareness is not reflected in the original treatment of knowledge about the existence and identity of objects with particular properties.

Referring expressions: Early on in the book, Appelt produces a most illuminating example to explain the need for planning at even the lowest levels of language generation.

> Two agents cooperating on a task [disassembling some complex piece of machinery] in which one has to make a request of the other. The speaker points to one of the tools on the table and says "Use the wheelpuller to remove the flywheel". The hearer...thinks to himself "Ah, so that's a wheelpuller. I was wondering how I was going to get that flywheel off." (p. 2)

This seems to me to be a marvelous illustration of the complexity and power of natural language. The only trouble is that nowhere in the remainder of the book do I see anything which gives me any real idea how to make it work. The nearest we see is an axiomatization of the act of describing with the following precondition axiom:

$$\begin{array}{l} \forall A,B,w1,w2\ R(:Do(A,\ :Describe(B,\ D)),\ w1,\ w2) \rightarrow \\ \quad \exists x\ T(w1,\ Know(@A,\ D(@x))\ \& \\ \quad T(w1,\ \neg\ Mutually\_know(@A,\ @B,\ \neg D^*(@x)))\ \& \\ \qquad T(w1,\ \forall y\ \neg\ Mutually\_know(@A,\ @B,\ D^*(y) \rightarrow y \neq @x)) \end{array}$$

Translated into English, this says that for A to use the referring expression D successfully as a description when talking to B, the following conditions must all hold. There must be an object which A knows fits the expression, A and B's mutual knowledge permits the description to fit the relevant object, and their mutual knowledge does not allow anything else to fit the description. There are, as with Appelt's other axiomatizations, several attractive things about this, especially that it permits each of them to know privately about other things which the description fits so long as they do not mutually know about them. Unfortunately there are again problems as well. Firstly, it fails to address the issue raised in the quote: what happens when A wants to use a term (wheel-puller) which B does not know the meaning of? It is clear that A can do this, and can even do it in the full knowledge that B does not yet know the meaning of the term in order to *convey* the meaning to him. Secondly, the assumption that we can compute what is *not* in A and B's kernel (which is what I take $\neg$Mutually_know(@A, @B, P) to denote) seems even less plausible than the assumption that we could compute what *is* in it. It is generally harder to show the nonexistence of a proof for something which is untrue than to show the existence of a proof for something which is true. Finally, the terms D and D* (which extends D by including facts about the discourse as well as about the object itself) are defined to be conjunctions of λ-expressions. Appelt therefore appears to be doing inference within a system that encompasses classical logic, modal logic and the λ-calculus. The history of mathematical logic since Frege and Russell is largely concerned with attempts to develop a consistent formalism within which this is possible. Some progress may currently being made via recent work on properties, but is certainly not the case that you can just plug λ-calculus into existing theorem provers and get any coherent results.

These remarks on things which KAMP seems unable to do should not be taken as condemnations. It is almost impossible for us to appreciate how difficult this task is, and *any* attempt is going to have problems until our understanding improves considerably. Appelt has done us all a favor by doing a good job of integrating a number of existing theories in an attempt to get a better overall picture of how language generation works. The value of the work resides in the effort that he has put into getting it all right. If this were a shoddy implementation of the various theories, the only lesson we could learn is that shoddy implementations are not good enough. But it does not feel like a shoddy implementation. The axiomatizations have been constructed with considerable care, and they do capture a number of significant properties of linguistic actions. If we feel at the end that it is still not good enough, the conclusion that we draw is that the entire approach is just not powerful enough yet. The reasons for trying to take some of the detail of the system apart and examine its limitations, then, are that since this is as good an implementation as we are likely to see of a language processing system based on planning theory and modal logic, any flaws we find will tell us what sort of a basis we need for the next step. Unfortunately it seems that this basis will need, at the very least, to include a treatment of intensions and a much more tentative approach to mutual knowledge. Language is more complicated than we thought.

One final note: my copy of this book lacked an index. This may or may not have been an accident with the copies sent out for review. It certainly made reading and cross-checking much of the detailed argument harder than would otherwise have been the case.

REFERENCES

Barwise, J., & Perry, J. (1983). *Situations and attitudes*. Cambridge, MA: Bradford Books.

Garrod, S., & Anderson, A. (1987). Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. *Cognition, 27*, 181–218.

Groenendijk, J., & Stokhof, M. (1987). *Type-shifting rules and the semantics of interrogatives*. Amsterdam: Institute for Language, Logic and Information, University of Amsterdam.

Grosz, B.J., Joshi, A., & Weinstein, S. (1983). Providing a unified account of definite noun phrases in discourse. *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics* (pp. 44-50).
Kay, M. (1985). Parsing in functional unification grammar. In D.R. Dowty, L. Karttunen, & A.M. Zwicky (Eds.), *Natural language parsing.* Cambridge, England: Cambridge University Press, 251-278.
Moore, R.C. (1980). *Reasoning about knowledge and action.* Menlo Park, CA: Stanford Research Institute Artificial Intelligence Center Report 191.
Moore, R.C. (1984). A formal theory of knowledge and action. In J. Hobbs & R.C. Moore (Eds.), *Formal theories of the commonsense world.* Norwood, NJ: Ablex.
Sacerdoti, E. (1977). *A structure for plans and behaviour.* Amsterdam: North Holland.
Searle, J. (1969). *Speech acts: An essay in the philosophy of language.* Cambridge, England: Cambridge University Press.
Sidner, C.L. (1983). Focussing in the comprehension of definite anaphora. In R. Berwick & J.M. Brady (Eds.), *Computational models of discourse* (pp. 331-371). Cambridge: MIT Press.
Turner, R. (1987). A theory of properties. *Journal of Symbolic Logic, 52,* 455-472.
Wallen, L. (1987). Matrix proof methods for modal logics. *International Joint Conference on Artificial Intelligence 87,* 917-923.

Arbib, M., Conklin, E., & Hill, J. ***From Schema Theory to Language.*** **New York: Oxford University Press, 1987. $39.95**

Reviewed by *Suzanne Mannes, University of Colorado, Boulder.*

In their book *From Schema Theory to Language,* Arbib, Conklin, and Hill have made a bold attempt at synthesizing several areas of cognitive science: biological control theory, neural modeling, artificial intelligence, cognitive psychology and brain theory. Their book includes the embellishment of several AI models which are representative of the current status of these areas. The three models discussed at length regard sentence comprehension, discourse generation about visual scenes and language acquisition in the child.

The main premise of the book is that we can view higher-level cognitive processes in a way similar to the way we understand visual processing as an action-perception cycle. We see something and our interpretation of that scene or object influences what we will see next. As we take in a visual scene we are constantly updating our representation of that scene and this in turn changes the future potential interpretations of that and other scenes. The world knowledge we have that allows us to make interpretations of relationships is stored in units called schemata, which share many features with semantic nets and production rules (p. 17). However, the chapter on knowledge representation is merely a rehashing of the well-known concepts of semantic nets and the differences between semantic networks and production rule formalisms and the relation between these representations and the