1

Architectures and Mechanisms in Sentence Comprehension

MARTIN J. PICKERING, CHARLES CLIFTON, JR., AND MATTHEW W. CROCKER

The architectures and mechanisms underlying language processing form one important part of the general architectures and mechanisms underlying cognition. In this book, the contributors largely focus on the question of what architectures and mechanisms underlie sentence comprehension, but it is our belief that their contributions illuminate the general nature of human language processing in the context of cognition as a whole. Because of the scope of these contributions, this introduction primarily concentrates on sentence comprehension. However, our perspective is to try to use evidence from sentence comprehension to understand the overall nature of language processing.

Let us first try to explain what we mean by architectures and mechanisms, and why these are important. We assume that the human language processor is a particular kind of computational device, and as such needs an organisation that we call its architecture. For instance, the processor might have two components, one dealing with word recognition, and one dealing with putting different words together into sentences. These processes might be completely distinct, in that the second processor could employ the output of the first processor, but have no impact on its internal workings. If so, we would know something about the architecture of the language processor. Architecture, therefore, refers to the static arrangement of components.

In contrast, the mechanisms are the dynamic workings of the language processor. A description of mechanisms is, in effect, a claim about the events in which the architectural components participate and about the causal structure of these events. If word recognition, say, works by decomposing words into component morphemes, processing each morpheme, and then putting the interpretations of the morphemes together, then this process would need a mechanism to describe it. A detailed understanding of the processes involved brings us closer to a description of the mechanisms that are used. Clearly, an understanding of the mechanisms relies on an understanding of the architecture, and vice versa.

The mechanism underlying some process, such as word recognition, can be described as being causally independent of other mechanisms only if a distinct component of the architecture deals uniquely with word recognition. If different architectural components are involved in the process of word recognition, then their causal interactions must be taken into account in describing the mechanisms underlying word recognition.

This description may give the impression that architecture is purely concerned with isolating components and fitting each component into the whole by showing which components impact on which other components. This is too simple. As Lewis (this volume) argues, a full description of cognitive architecture goes beyond "module geography" (J. D. Fodor, 1990, cited in Lewis). Lewis proposes that a specification of a functionally complete architecture must include (at least) a specification of the memories and the primitive processes the modules use, a specification of how the modules pass information among themselves, and a specification of how processes are initiated and terminated. This should include answers to questions such as what processes can go on in parallel and how processes must be arranged so that one must complete before the next can begin.

Why should understanding architecture and mechanisms be important? We might believe that the interesting things to understand about language concern the question of why we say what we say, why conversations work as they do, and so on. Of course these are interesting questions, but they are different from the issue at hand here. The question of understanding architectures and mechanisms is no different from the question of understanding the architectures and mechanisms that underlie other aspects of cognitive psychology, such as, for example, visual perception, skilled performance, or reasoning. Ultimately, we are concerned with a particular cognitive process, that of language processing, and with how it relates to the rest of cognition.

Let us now concentrate on the specific question of sentence comprehension. This is concerned with the way in which people determine the meaning of a sentence. Sentence comprehension lacks a simple definition, yet its general boundaries can be identified. Most obviously, it is concerned with the process of going from sound or written text to meaning. It assumes that words have been recognised. What it is principally concerned with is the way that the processor combines the meaning of words to produce a meaning for a sentence as a whole. A large part of this process involves parsing: determining how people obtain a particular syntactic analysis for a sentence or sentence fragment. This has been the traditional focus of sentence processing research and describes the majority of the contributions to this book. However, sentence comprehension goes beyond syntactic analysis. It includes aspects of semantic interpretation, as found in the processing of quantifiers like *every* and *a*. For example, *Every boy climbs a tree* is ambiguous between two interpretations, depending on whether each boy climbs the same tree or not. A theory of sentence comprehension needs to explain how the processor constructs and selects a particular interpretation. Likewise, pronouns like *he* or *she* are often ambiguous, but the processor normally makes a decision about which entity they refer to. In general, we can say that theories of sentence comprehension draw upon the assumptions and distinctions employed in theoretical linguistics, primarily in syntax and formal semantics (see Almor, this volume; Frazier, this volume; Moxey & Sanford, this volume).

General knowledge plays an extremely important role in sentence comprehension. Consider (1) and (2):

(1) The boy saw the girl with the flowers.

(2) John gave Bill some food because he was hungry.

Both of these sentences are formally ambiguous, but in both cases general knowledge disambiguates them. In (1), the prepositional phrase *with the flowers* can attach high, to the verb phrase *saw the girl*. But this produces the implausible meaning that the boy used the flowers to see the girl. Therefore, the prepositional phrase has to attach low, to the noun phrase *the girl*. This produces the plausible meaning that the boy saw the girl who had the flowers. In (2), the antecedent of *he* could be *John*. But this leads to the implausible interpretation that John gave away food because he was hungry, so the antecedent of *he* has to be *Bill* instead. In both cases, general knowledge is used to determine the appropriate analysis. Sentence comprehension is concerned with the impact of general knowledge on language comprehension.

However, the study of general knowledge itself is not part of psycholinguistics, and sentence comprehension is not interested in all processes that employ general knowledge. For instance, it is not primarily concerned with the issue of what inferences are drawn on the basis of the final interpretation assigned to a sentence. Resolving this question is presumably a goal of an account of general cognition, rather than of sentence comprehension in particular. In order to be able to make this distinction, we need to assume that sentence comprehension produces a kind of output that can be employed by other cognitive systems (cf. J. A. Fodor, 1983). This is a common, if not universal, assumption.

In this introductory chapter, we discuss the mechanisms that may be employed in sentence comprehension and discuss hypotheses about the architecture of the processor. The field of sentence comprehension is extremely controversial in many respects. Here, we attempt to identify points of agreement and to circumscribe the areas of debate. Researchers certainly disagree about the actual architectures and mechanisms, but they appear to agree about the questions that need to be addressed, about the range of reasonable positions, and about the types of methodologies that are appropriate.

Moreover, we believe that researchers may actually agree about actual architectures and mechanisms to a greater extent than at first appears. For example, all theories assume that sentence comprehension is highly incremental: once a word is recognised, it is rapidly integrated with preceding context. In addition, there is little dispute about the sources of information that are relevant to sentence comprehension. In areas of considerable controversy, there is often agreement about the kinds of evidence that could be used to distinguish theories. Most of the arguments concern broad issues of architecture and mechanism, such as whether the processor works in serial or in parallel, whether it always has to construct a fully-specified interpretation, when it is able to draw upon different sources of information, and whether it employs statistical information in choosing an analysis. Our introduction considers these issues, amongst others. In most cases, we employ examples from parsing, but many of these issues also apply to other aspects of sentence comprehension. First, we highlight the highly incremental nature of most aspects of sentence comprehension.

1.1 Incrementality

In the comprehension of both spoken and written language, words are encountered sequentially. Experimental evidence indicates that a great deal of processing occurs immediately, before the next word is encountered (e.g., Just & Carpenter, 1980; Tyler & Marslen-Wilson, 1977). A very important effect of this is that decisions are often made when a sentence fragment is ambiguous in some way. Word recognition is not normally delayed (e.g., Marslen-Wilson, 1987; Rayner & Duffy, 1986), even if disambiguation occurs after the ambiguous word (e.g., Rayner & Frazier, 1989). More relevantly, there is normally no measurable delay before syntactic analysis and some aspects of semantic interpretation begin (though it is impossible to be sure that there are no circumstances under which delay occurs). However, some higher level aspects of processing may be delayed – it is an open question how far incrementality goes.

Evidence for incremental parsing comes from the vast literature showing garden-path effects. Briefly, the sentence *The horse raced past the barn fell* (Bever, 1970) is hard (in part, at least) because people initially assume that *raced* is an active past-tense verb, and hence that *the horse raced past the barn* is a complete sentence. When they encounter *fell*, they realise this is impossible and reinterpret *raced* as a past participle in a reduced-relative construction (cf. *The horse that was raced past the barn fell*), or fail to understand the sentence entirely. In other words, they are "led up the garden path" by such a sentence. From this, we can tell that they have performed incremental syntactic analysis and committed to it, before reaching *fell*. Experimental evidence strongly supports this conclusion for many different sentence types and suggests that syntactic analysis begins very rapidly (e.g., Frazier & Rayner, 1982).

Evidence also supports the intuition that people start to understand sentences as they hear or read them. Most famously, Marslen-Wilson (1973, 1975) showed that subjects' errors in shadowing a text at a lag of only 300 milliseconds were constrained by semantic context. This demonstrates that the meaning of what is heard can be rapidly integrated with general knowledge (though it is conceivable that integration occurs during production rather than comprehension). Data from eye-tracking gives more direct evidence for incremental interpretation. For instance, Traxler and Pickering (1996b) found that readers were disrupted as soon as they read the word *shot* in (3):

(3) That is the very small pistol in which the heartless killer shot the hapless man yesterday afternoon.

Hence, they must have semantically processed the sentence fragment up to *shot* when they first encounter the word. Many other experiments using various methods also provide good evidence for incremental semantic processing (e.g., Boland, Tanenhaus, Garnsey, & Carlson, 1995; Clifton, 1993; Garrod, Freudenthal, & Boyle, 1994; Holmes, Stowe, & Cupples, 1989; Swinney, 1979; Trueswell, Tanenhaus, & Garnsey, 1994; Tyler & Marslen-Wilson, 1977). Hence, the language processing system must very rapidly construct a syntactic analysis for a sentence fragment, assign it a semantic interpretation, and make at least some attempt to relate this interpretation to general knowledge.

There may, however, be some limits to incrementality. For instance, Frazier and Rayner (1987) claimed that syntactic category ambiguities may not be resolved immediately (e.g., in *the desert trains* the word *trains* is not immediately categorised as a noun or a verb); see Corley and Crocker (this volume) for discussion. In a similar vein, Frazier and Rayner (1990) suggested that readers may not resolve the distinct senses of words like *newspaper* (the concrete object vs. the abstract institution sense). However, incrementality does seem to be the norm, even when it might not be expected. For example, a case where syntactic analysis might reasonably be delayed occurs in verb-final languages, such as Dutch and Japanese, where the parser often encounters more than one noun phrase before a verb. Evidence suggests that the parser does attach and interpret these noun phrases as co-arguments, before the subcategorising verb is reached (e.g., Bader & Lasser, 1994; Crocker, 1994; Frazier, 1987b; Sturt & Crocker, 1996). This rules out, or at least entails substantial revision of, head-driven and delay models of parsing such as Marcus (1980), Abney (1989) and Pritchett (1992). Another possible case of delay relates to "non-primary" phrases (Frazier & Clifton, 1996).

In anaphoric resolution, the evidence for incrementality is more controversial, and it is possible that only some anaphoric resolution is immediate. Good evidence for incrementality comes from Marslen-Wilson, Tyler, and Koster (1993) in listening and from Garrod et al. (1994) in reading. However, Garrod et al.'s study suggested that pronouns are only resolved immediately if they refer to a well-focused and unambiguous antecedent. In addition, Greene, McKoon, and Ratcliff (1992) presented data from a priming task that suggests that readers may not, under some conditions, bother to resolve the antecedent of an unambiguous pronoun immediately. Another interesting area concerns the resolution of quantifier-scope ambiguities, but here there is little evidence about incremental processing (e.g., Kurtzman & MacDonald, 1993 investigated preferred interpretations, but only by monitoring reading time on the next sentence).

Finally, it is possible that incrementality might be clause-bound, and that the processor links the interpretation of different clauses at clause boundaries. For instance, Millis and Just (1994) argued that readers process two clauses linked by a connective like *because* by interpreting each clause separately and only combining them at the end of the second clause. However, Traxler, Bybee, and Pickering (1997) found evidence that the processor detected implausible connections between clauses before the end of the second clause, so there is at least some evidence that incrementality extends to higher levels of semantic interpretation.

We conclude that a major constraint on the architectures and mechanisms underlying language processing is that they must support extremely incremental processing during comprehension. One important consequence is that the processor must be able to employ a great deal of information very rapidly. For instance, in syntactic processing it must be able to determine the syntactic category of a word when it is encountered. To do this, it needs to have access to the relevant syntactic information and to be able to apply this knowledge to the specific task that it faces. We now need to consider what these knowledge sources might be, and how the processor employs this information.

1.2 Knowledge Sources

In the 1980s, discussions of language processing often focused on the general question of whether the language processor was a module (e.g., J. A. Fodor, 1983; Forster, 1979; Rayner, Carlson, & Frazier, 1983; Taraban & McClelland, 1988; Tyler & Marslen-Wilson, 1987; see Garfield, 1987). But more recently, many researchers have started to focus on when and how specific sources of information are employed.

Most researchers assume representational modularity (Trueswell et al., 1994): sources of information like syntactic and semantic knowledge are represented separately. This assumption is standard to most generative grammar (e.g., Chomsky, 1965, 1981; Pollard & Sag, 1993; but see, e.g., Lakoff, 1986), and is assumed in the great majority of psycholinguistic research (though see McClelland, St. John, & Taraban, 1989). Some experimental evidence provides support for the standard position. For example, it is possible to prime syntactic structure in a manner probably independent of semantic factors, both in production (Bock, 1986) and comprehension (see Branigan, Pickering, Liversedge, Stewart, & Urbach, 1995). Also, evidence from event-related brain potentials (where electrical activity from the brain is measured as subjects perform a task) suggests that syntactic and semantic processing are distinct. Semantically anomalous words elicit a negative-going brain wave about 400 ms after the stimulus (Kutas & Hillyard, 1980; see also Garnsey, Tanenhaus, & Chapman, 1989). For instance, after *I drink my coffee with cream and*, the wave for the anomalous word *dog* is negative compared with the wave for *sugar*. In contrast, syntactically anomalous words produce a positive-going wave around 600 ms after the stimulus, whether the anomaly is due to ungrammaticality or to a garden-path construction (Hagoort, Brown, & Groothusen, 1993; Osterhout & Holcomb, 1992; Osterhout, Holcomb, & Swinney, 1994; Osterhout & Mobley, 1995; see Brown & Hagoort, this volume).

Assuming representational modularity, relevant sources of information include:

I. Syntactic Category Information

The lexical entry for each word indicates its syntactic properties. For example, the entry for *vanishes* states that it is an verb and that it is intransitive (taking a subject but no object). Psycholinguistically, this might count as one source of information or as two, if major-category information (being a verb) is separated from subcategorisation information (being intransitive). If the latter is the case, then the processor might base initial processing decisions on major-category information alone (e.g., Ferreira & Henderson, 1990; Mitchell, 1987). Also,