1 The Case for Interactionism in Language Processing

James L. McClelland
Carnegie-Mellon University
Pittsburgh, Pennsylvania, U.S.A.

ABSTRACT

Interactive models of language processing assume that information flows both bottom-up and top-down, so that the representations formed at each level may be influenced by higher as well as lower levels. I describe a framework called the *interactive activation* framework that embeds this key assumption among others, including the assumption that influences from different sources are combined nonlinearly. This nonlinearity means that information that may be decisive under some circumstances may have little or no effect under other conditions. Two attempts to rule out an interactive account in favour of models in which individual components of the language processing system act autonomously are considered in the light of the interactive activation framework. In both cases, the facts are as expected from the principles of interactive activation. In general, existing facts do not rule out an interactive account, but they do not require one either. To demonstrate that more definitive tests of interaction are possible, I describe an experiment that demonstrates a new kind of influence of a higher-level factor (lexical membership) on a lower level of processing (phoneme identification). The experiment illustrates one reason why feedback from higher levels is computationally desirable; it allows lower levels to be tuned by contextual factors so that they can supply more accurate information to higher levels.

INTRODUCTION

When we process language—either in written or in spoken form—we construct representations of what we are processing at many different levels. This process is profoundly affected by contextual information. For example, in reading, we perceive letters better when they occur in words. We recognise words better when they occur in sentences. We interpret the meanings of

words in accordance with the contexts in which they occur. We assign grammatical structures to sentences, based on the thematic constraints among the constituents of the sentences. Many authors—Huey (1968), Neisser (1967), and Rumelhart (1977), to name a few—have documented some or all of these points.

Clearly, this use of contextual information is based on what we know about our language and about the world we use language to tell each other about. How does this knowledge enter into language processing? How does it allow contextual factors to influence the course of processing?

In this paper, I will describe a set of theoretical principles about the nature of the mechanisms of language processing that provides one possible set of answers to these questions. These principles combine to form a framework which I will call the *interactive activation* framework. The paper has three main parts. In the first part, I will describe the principles and explore a central reason why they offer an appealing account of the role of knowledge in language processing. In the second part, I will consider two prominent lines of empirical investigation that have been offered as evidence against the view that particular parts of the processing system are influenced by multiple sources of information, as the interactive activation framework assumes. Finally, in the third part, I will discuss one way in which interactive processing might distinguish itself empirically from mechanisms that employ a one-way flow of information.

To summarise the main points of each part:

1. In the interactive activation framework, the knowledge that guides processing is stored in the connections between units on the same and adjacent levels. The processing units they connect may receive input from a number of different sources. This allows the knowledge that guides processing to be completely local, while at the same time allowing the results of processing at one level to influence processing at other levels, both above and below. Thus, the approach combines a desirable computational characteristic of an encapsulationist position (Fodor, 1983) while retaining the capacity to exploit the benefits of interactive processing.

2. Two sources of empirical evidence that have been taken as counting against interactionism do not stand up to scrutiny. The first case is the resolution of lexical ambiguity in context. Here I re-examine existing data and compare them with simulation results illustrating general characteristics of interactive activation mechanisms to show that the findings are completely consistent with an interactive position. The second case considered is the role of semantic constraints in the resolution of syntactic ambiguities. Here I review some recent data that demonstrate the importance of semantic factors in phenomena that had been taken as evidence of a syntactic processing strategy that is impervious to semantic influences. In both cases I will argue

that the evidence is just what would be expected from an interactive activation account.

3. It is an important and challenging task to find experimental tests that can distinguish between an interactive system and one in which information flows only in one direction. Unidirectional and interactionist models can make identical predictions for a large number of experiments, as long as it is assumed that lower levels are free to pass on ambiguities they cannot resolve to higher levels. However, experimental tests can be constructed using higher-level influences to trigger effects assumed to be based on processing at lower levels. I will illustrate this method by describing a recent experiment that uses it to provide evidence of lexical effects on phonetic processing, and I will suggest that this method may also help us to examine higher-level influences on lower levels of processing in other cases.

THE INTERACTIVE ACTIVATION FRAMEWORK

The following principles characterise the interactive activation framework. These principles have emerged from work with the interactive activation model of visual word recognition (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982), the TRACE model of speech perception (Elman & McClelland, 1986; McClelland & Elman, 1986), and the programmable blackboard model of reading (McClelland, 1985; 1986). The principles apply, I believe, to the processing of both spoken and written language, as well as to the processing of other kinds of perceptual inputs; however, all the examples I will use here are taken from language processing.

The Processing System is Organised Into Levels.   This principle is shared by virtually all models of language processing. Exactly what the levels are, of course, is far from clear, but this is not our present concern. For present purposes, I will adopt an illustrative set of levels to provide a context in which to discuss the processing interactions that may be involved in reading a sentence. These levels are a visual feature level, a letter level, a word level, a syntactic level, a word-sense level, and a scenario level, on which the representation captures the nonlinguistic state or action described by the sentence being processed. Higher levels are, of course, required for longer passages of text, but the set of levels will provide a sufficient basis for the phenomena we will consider here. For processing speech, we also need a phonetic level and an auditory feature level to provide input to the phonological level.

The Representation Constructed at Each Level is a Pattern of Activation Over an Ensemble of Simple Processing Units.   This assumption is central to

the entire interactive activation approach, and strongly differentiates it from other approaches. In this approach, representations are active—they can influence, and be influenced by, representations at other levels of processing. In this paper, I will adopt the formal convenience of assuming that individual processing units stand for individual conceptual objects such as letters, words, phonemes, or syntactic attachments. Thus, a representation of a spoken word at the phonetic level is a pattern of activation over units that stand for phonemes; these units are role-specific, so that the pattern of activation of "cat" is different from the pattern of activation of "tac."

Activation Occurs Through Processing Interactions that are Bi-directional, Both Within Levels and Between Levels. A basic assumption of the framework is that processing interactions are always reciprocal; it is this bi-directional characteristic that makes the system interactive. Bi-directional excitatory interactions between levels allow mutual simultaneous constraint among adjacent levels, and bi-directional inhibitory interactions within a level allow for competition among mutually incompatible interpretations of a portion of an input. The between-level excitatory interactions are captured in these models in two-way excitatory connections between mutually compatible processing units; thus the unit for word-inital /t/ has an excitatory connection to the unit for the word /tac/, and receives an excitatory connection from the unit for the word /tac/.

Between-level Processing Interactions Occur Between Adjacent Levels Only. This assumption is actually rather a vague one, since adjacency itself is a matter of assumption. I mention it because it restricts the *direct* processing interactions to a reasonably small and manageable set, rather than allowing everything to influence everything else directly. One possible set of interactions between levels is sketched in Fig. 1.1. Note that even though some pairs of levels are not directly connected, each level can influence each other level indirectly, via indirect connections.

Between-level Interactions are Excitatory Only; Within-level Interactions are Competitive. A feature of the interactive activation framework that has gradually emerged over the years is the idea that between-level interactions should be excitatory only, so that a pattern of activation on one level will tend to excite compatible patterns at adjacent levels, but will not directly inhibit incompatible patterns. The inhibition of incompatible patterns is assumed to occur via competition among alternative patterns of activation on the same level. This idea is characteristic of assumptions made by Grossberg (1976 and elsewhere), and its utility has become clearer in later versions of interactive activation models (McClelland & Elman, 1986; McClelland, 1985). The principal reason for this assumption is that it allows
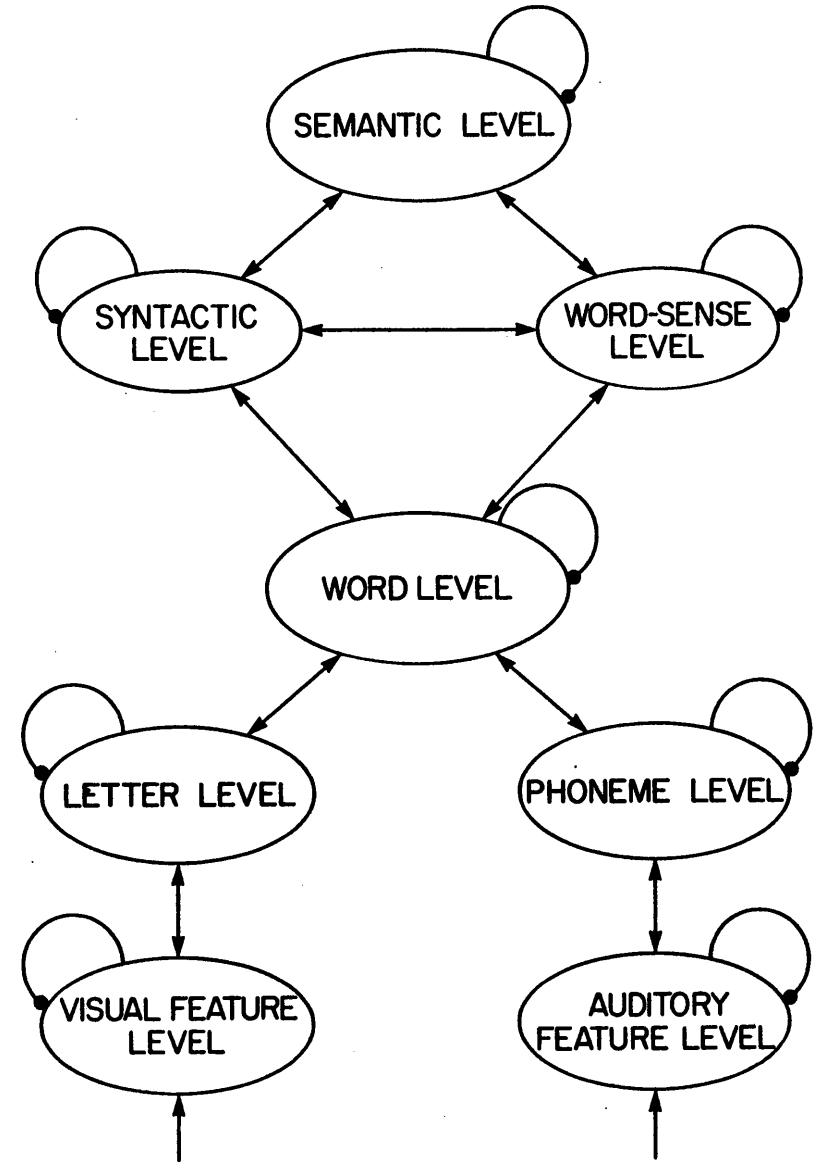
FIG. 1.1.   A set of possible processing levels and connections among these levels. In an interactive activation model, each level would consist of a large number of simple processing units. No claim is made that this is exactly the right set of levels; this set is given for illustrative purposes only. Bi-directional, excitatory connections are represented by double-headed arrows between neighbouring levels. Inhibitory within-level connections are represented by the lines ending in dots that loop back onto each level.

possible alternative representations to accumulate support from a number of sources, then to compete with other alternative possibilities so that the one with the most support can dominate all the others. This allows the network to implement a "best match" strategy of choosing representations; for example, a sequence of phonemes that does not exactly match any particular word will nevertheless activate the closet word. Thus "parageet" for example can result in the recognition of the word "parakeet," even though it does not match parakeet exactly.

Activations and Connections are Continuously Graded.   The activation of a representation is a matter of degree, as is the strength of the influence one a representation exerts on another. Degree of activation of a unit reflects the strength of the hypothesis that the representational object the unit stands for is present; the strengths of the connections between units reflect the strengths of the contingencies that hold between the representational objects.

The Activation Process is Nonlinear.   Each processing unit in an interactive activation network performs a very simple computation. It adds up all of the weighted excitatory influences it receives from other units and subtracts from these the weighted inhibitory influences that it receives from competing units. Then, it updates its activation to reflect this combined (what I will call *net*) input. The activation of the unit is monotonically, but not linearly, related to this sum; at high levels of excitatory input, activation levels off at a maximum value, and with strong inhibitory input, it levels off at a minimum value. Because of these nonlinearities, and because of the competitive interactions among units, inputs that are sometimes crucial for determining the outcome of processing may have little or no effect at other times.[1] The specific details of the nonlinear activation assumptions that I have used are based on, though not identical with, those used by Grossberg (e.g., Grossberg, 1978).

Activation Builds Up and Decays Over Time.   It is assumed that processing interactions occur continually, but that the activation process is gradual and incremental, so that it takes time for activation to propagate through the system. New inputs begin to have their effects immediately, but these effects build up over time and then gradually decay away as processing continues.

---

[1]It is worth noting that this nonlinear characteristic is absolutely essential to the operation of the network as a whole; if all units in the system behaved linearly, no purpose would be served by having multiple levels, and none but the most trivial of computational operations could be performed. Furthermore, feedback from higher levels to lower levels can lead to runaway activation in a linear system. For discussion, see Rumelhart, Hinton, & McClelland (1986).

These assumptions are now being applied in the construction of models of higher-level aspects of language processing, such as the assignment of constituents of sentences to semantic roles and disambiguation of word meaning in context (Cottrell, 1985; Waltz & Pollack, 1985; Kawamoto, Note 4; McClelland & Kawamoto, 1986). At higher levels of processing, I and other researchers have tended to build models that make explicit use of distributed representation, in which a conceptual object is represented by a pattern of activation, rather than a single unit (Hinton, McClelland & Rumelhart, 1986). However, even here it is convenient to speak of whole patterns of activation as though they were separate information-processing constructs, that interact with each other via excitatory and inhibitory contingencies. Indeed the distributed representation can be seen as an implementation of the more abstract, functional description (see Smolensky, 1986 for a discussion of this issue).

Encapsulated Knowledge, Interactive Processing

In his book on modularity, Fodor (1983) explains a virtue of dividing up the knowledge that is used, and encapsulating portions of it in separate modules each dedicated to a specific part of a complex information processing task. Encapsulation of knowledge allows, he notes, for automatised, reflex-like processing in each module, since each module need only consult a finite store of locally-relevant information.

The interactive activation framework adheres to this desirable property. A central feature of the framework is the fact that the knowledge that guides processing is intrinsically local and inaccessible to other portions of the network. To see this, it is useful to focus attention on the connections between some pair of adjacent levels in the system; for example, the connections from the letter level to the word level. These connections are the knowledge that allows the system to form appropriate word level representations from patterns of activation at the letter level. They express contingencies between activations of units at the letter level, and activations of units at the word level. This information is completely encapsulated within this part of the processing mechanism; it is never consulted by any other part of the mechanism. By the same token, this part of the mechanism never consults the knowledge stored in any other part in doing its job, which is simply to supply input to the units at the word level. We have, then, a system in which the knowledge is completely encapsulated.

At the same time, the architecture of the system overcomes what I believe is an unnecessary limitation that Fodor places on modular systems; that is that the output of a module be independent of influences from other sources. Interactive activation provides a framework for processing in which multiple sources of information can influence the construction of representations at

each level. This is because each level combines inputs it receives from multiple sources in determining what its pattern of activation shall be. The input a level receives from a particular adjacent level, then, simply constitutes one source of constraint on the construction of a representation that is subject to influence by other sources.

Where Fodor's analysis went astray, I believe, is in assuming that the combined use of constraints from multiple sources requires each module in the system to have access to knowledge of many different types. What the interactive activation framework makes clear is that this is not the case. Each processing level—each set of units—provides a device that performs a very general computation that allows it to combine inputs from a number of sources. This general computational characteristic of interactive activation mechanisms provides a simple way for knowledge at all different levels to exert simultaneous influence on the outcome of processing, without requiring any part of the system to know very much at all.[2]

AN EXAMINATION OF THE EVIDENCE

No-one doubts that the ultimate outcome of processing is sensitive to influences from many levels. The psychological literature is replete with demonstrations of such effects; but many researchers have questioned the view that the influences exerted by higher levels occur through direct influences from higher levels back down into lower levels of processing. There are two poles to this argument. First, the results of some experiments have been taken as evidence against an interactive view, at least with respect to certain aspects of processing. Second, it is often pointed out that results that could be attributable to interactive processing might be explained in other ways; Fodor (1983) makes this point repeatedly.

I will consider these two aspects of the argument against interactionism in turn. First I will consider two cases of experimental findings that have been taken as evidence against interactionism in two specific cases. Here my aim is to show that the experimental facts, when looked at closely, turn out to be perfectly consistent with an interactive activation account. I do not mean to say that they cannot be interpreted without recourse to interaction between levels. Though the phenomena are just what we expect from an interactive

[2]I should note that Fodor suggests reasons other than computational efficiency for advocating autonomy of processing. For one thing, he suggests that if modules are autonomous it may be easier for cognitive scientists to analyse exactly what functions each module computes. While this might well be the case, it seems unlikely that the convenience of cognitive scientists entered into the design of our computational machinery; computational considerations seem more likely to have influenced the course of evolution; and my argument is that such considerations favour interactionism.

activation approach, there can be alternative interpretations. In a later section, I will turn specifically to the question of how one might find evidence that more clearly favours an interactive activation view.

The Case Against Interactionism

The two cases I will consider both purport to demonstrate the autonomy of some aspect of processing from higher-level, or contextual influences. One of these cases concerns accessing word meanings. The other concerns the mechanism that determines how constituents should be attached to each other in constructing a representation of the syntactic structure of a sentence.

In examining each of these cases, it will be helpful to have two basic properties of interactive systems in view. The first is that contextual influences often produce what I will call selective, as opposed to predictive, effects. The second is that contextual effects—indeed, the effects of any factor—can be masked by strong effects of other factors. The first fact will be useful when we come to interpret evidence that context appears to exert primarily a selective effect in certain lexical ambiguity resolution experiments; the second will be most relevant when we examine evidence that semantic context effects do not show up in the initial processing of certain grammatical constructions.

To illustrate the first point, let us consider the recognition of an ambiguous phoneme embedded in a context which should favour one interpretation over the other. A simulation illustrating this is shown in Fig. 1.2, using the TRACE model of speech perception (McClelland & Elman, 1986).

To understand the simulation, some facts about the model are necessary. The model consists of units grouped into three processing levels. There is a phonetic feature level, a phoneme level, and a word level. Within each level, there are separate pools of units for each small temporal segment of an utterance. Thus successive phonemes in a word activate phoneme detectors in successive pools of units. It is useful to visualise the feature units as though they are laid out in successive banks from left to right in space, with banks of phoneme units above them and banks of word units above the phoneme units. Each bank of units covers only a small temporal window. Spoken input is swept across this spatial array from left to right, providing input to feature units in successive banks as time progresses. Connections between feature and phoneme units allow active feature units in a particular bank to send excitatory input to units for appropriate phonemes in corresponding banks; phoneme-to-word connections allow phonemes to send excitation to appropriate words in corresponding banks; there are also feedback connections from the word level to the phoneme level and from the phoneme to the feature level. In addition to these excitatory connections, there are also inhibitory connections between units which span overlapping temporal
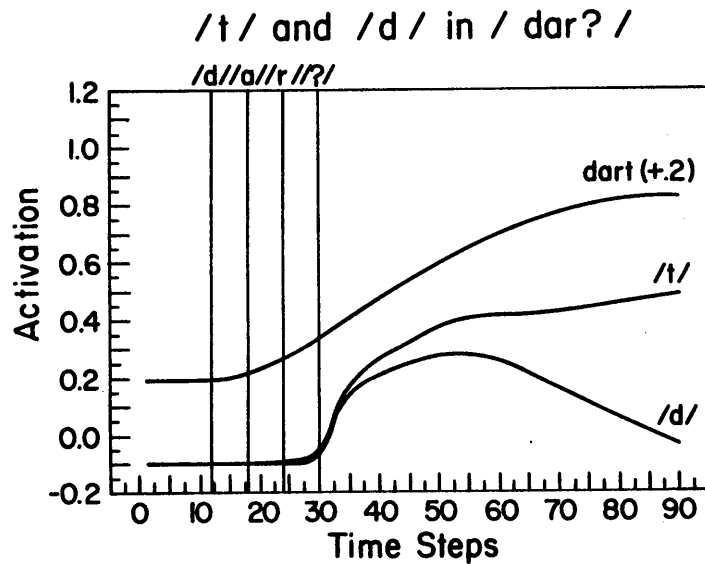
/t/ and /d/ in / dar? /

FIG. 1.2.   The time course of activation of units for /d/ and /t/ at the end of the string /dar?/, where the ? stands for a segment ambiguous between /t/ and /d/. The time course of activation of the unit for the word *dart* is also shown.

regions. At the phoneme level, this means that competition occurs only among alternative phonetic interpretations of the same temporal segment of speech.

In our example, we will consider an input that consists of the phonemes /d/, /a/, and /r/, followed by a phonetic segment that is ambiguous between /d/ and /t/. The figure illustrates the build-up of activation for the phoneme units activated by the final ambiguous sound. We can see that, initially, there is a very slight advantage of the /t/ over the /d/. This advantage stays relatively constant for a time, but gradually /t/ begins to dominate /d/ and to push its activation down. While both phonemes are activated initially, only one remains active in the end.

Why is the context effect so small at first? The primary reason has to do with the degree of constraint imposed by the context. Activation of the /t/ over the /d/ results from feedback from the word level, but at the time the /t/ and /d/ are coming in, the relevant word detector (for the word *dart*) is not very active. The reason is simply that there are several other words that are still consistent with the input up to that point. These words are all in competition, so that none are very highly activated. The ambiguous phoneme itself must determine which of these words is really being said, and thereby

allow it to dominate the possibilities left open by preceding portions of the input. Only after the ambiguous word strengthens the activation of *dart* over its competitors can *dart* really provide strong support for the /t/ interpretation of the final phoneme.

I want to make it clear that context can and does exert stronger effects than we see here under some circumstances. When, for example, an ambiguous segment comes at the end of a long word that has no remaining competitors a few phonemes before the ambiguous segment is received, we see much stronger context effects in the simulation. These effects are, of course, consistent with the empirical finding that lexical effects in speech processing are larger at later points in words (Marslen-Wilson & Welsh, 1978; Samuel, 1981).

The essential point is that context that is clearly strong enough to exert a potent role in determining the eventual outcome of processing may very well exert its influence primarily by selecting among alternatives as they are becoming activated bottom-up. An initial slight advantage is generally observed for the contextually appropriate alternative, but both appropriate and inappropriate alternatives may receive considerable activation before the resolution of the ambiguity is complete.

Now we consider the second point, namely that effects of context can be blocked if there are other factors that are exerting stronger influences. To demonstrate this, I will show the results of two more simulation runs with the TRACE model, using an unambiguous final /d/ in one case and an unambiguous final /t/ in the other, preceded by the string /dar/. Here context should support the /t/, since *dart* is a word. However, as Fig. 1.3 shows, when the input is unambiguous it produces strong bottom-up support for the phoneme actually presented, and this actually blocks out the effect of context almost completely.

Though there is a slight advantage for the /t/, it is very small and might easily go undetected in an experiment. Certainly, there is no doubt that a /t/ will be heard in one case and a /d/ in the other. The reason is that with strong bottom-up input favouring a particular interpretation, the correct answer is quickly locked into the system and keeps the alternatives from becoming activated, due to competitive inhibition among units standing for alternative interpretations at the same level. The differential feedback support that the /t/ receives does not really become strong enough to influence processing until it is too late.

Again, I want to make clear that the effect of context would be stronger in other cases. When there is a strong expectation before the target occurs, feedback from higher levels can act as a second source of excitation favouring the one alternative; under these conditions, the contextually favoured alternative will have more of an advantage. But in many cases, a context that would be sufficient to disambiguate a borderline stimulus, as we

## /d / in / dard /and / t /in / dart /

FIG. 1.3.   Time course of activation of detectors for the final /t/ in /dart/ and the final /d/ in /dard/. Also shown is the time course of activation of the detector for the word *dart* in each case.

saw in the previous simulation, will have very little effect when the stimulus is not borderline, as in the present case.

These kinds of effects, where a strong cue overshadows the effects of a weak cue that is known to operate under other circumstances, are absolutely ubiquitous in the literature. They are nicely explained by the interactive activation approach, and by other models such as the Oden-Massaro information integration model (Oden & Massaro, 1978). As just one example, Ganong (1980) found just these kinds of effects in his initial studies of the lexical effect in phoneme identification. He reported that context biased the interpretation of ambiguous sounds at or near the boundary between two phonetic categories, but did not alter the interpretation of unambiguous sounds well within one category or another. One hears the /k/ in (strongly articulated) *kift* correctly, in spite of the unfavourable context. Simulations reported in Elman and McClelland (1986) show that these sorts of effects are expected in the interactive activation framework.

Given these preliminary observations, we are now ready to consider the case against interaction in lexical access and in syntactic analysis. In the first case, the claim has been made that initial access to words occurs autono-

mously, without regard to context, and that higher levels simply select the appropriate word from those that are made available by the autonomous access mechanism (Tanenhaus, Leiman, & Seidenberg, 1979; Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982). In the second case, the claim is that the syntactic processing of a sentence is encapsulated, so higher levels of processing only accept or reject possible parses presented to them by the syntactic level. I've chosen to examine these cases for two reasons. First, they are both often cited as evidence of autonomy, and so they are worth considering, in and of themselves. Second, they each illustrate characteristics of the interactive activation framework that ought to be taken into account in attempts to argue against an interactive position.

Word Sense Disambiguation

There are now several studies using a cross-modality priming paradigm to study word sense disambiguation. The first two such studies were those of Tanenhaus et al. (1979) and of Swinney (1979). In these and other studies, the following pattern has been found: Immediately after an ambiguous word, both meanings appear to be activated, even when context is provided which favours one interpretation of the target word over the other. After a delay, the only contextually appropriate meaning appears to remain active.

This pattern of results has been interpreted as favouring a view that I will call the *autonomous lexical access position* (Tanenhaus et al., 1979). According to this position, the process of accessing meanings of words is driven only by the bottom-up processing of the stimulus; context operates only later, to select among the alternatives that are made available by the bottom-up access process.

In this section, I will argue that the results indicate instead a pattern that conforms to what we would expect from an interactive activation model: Initially both meanings appear to be accessed, but—and this is the crucial point—the evidence suggests that the contextually appropriate reading is in fact favoured over the contextually inappropriate reading, even early on in processing.

In documenting this claim, I will focus first on the experiments of Swinney (1979). He presented ambiguous words like "bugs" in contexts which favoured one or the other meaning of this word (insects or snooping devices). The ambiguous word occurred in a spoken passage, and subjects listened to the passages through earphones; at the end of the ambiguous word, they were tested with a visually presented probe word. This word could be related to the contextually appropriate meaning of the ambiguous prime word (*ants*), to the contextually inappropriate meaning (*spy*), or it could be unrelated to the ambiguous word (*sew*). The task was simply to indicate whether the

visually presented probe was a word or not. Nonword probes were of course presented on other trials.

The results of Swinney's experiment showed faster lexical decision reaction times to probes related to both meanings of the ambiguous prime word, relative to control. There was a 70msec advantage for the target related to the contextually appropriate meaning of the ambiguous prime, and a 50msec advantage for the target related to the contextually inappropriate meaning of the prime. Both were significantly faster than the responses in the control condition.

In a follow-up study, Swinney replicated his first experiment, and compared the results to the results of a second condition, in which the probe was delayed by three syllables. At 0 delay, the appropriate probe showed 38msec facilitation and the inappropriate probe showed 31msec. After the delay, the appropriate probe showed 47msec and the inappropriate probe was 1msec slower than control. Because the second experiment contains all of the relevant conditions, I have graphed the results in Fig. 1.4.

The basic pattern of results obtained by Swinney was also found by Tanenhaus et al. (1979), hereafter called TLS, and by Seidenberg et al. (1982), hereafter called STLB. In fact, in two conditions of STLB (for noun-noun ambiguities in Experiments 2 and 4) there was a significant selective priming effect at 0 delay. However, in four other conditions over the two experiments, priming of both meanings was found. Looking just at the six different experiments finding priming of both meanings at 0 delay (two of Swinney's, one from TLS, and three from STLB) we find that in five of the six cases, the contextually appropriate target receives stronger priming than the inappropriate one. These findings are summarised in Table 1.1. TLS and STLB also provide confirmation that at a delay, there is strong selection of the contextually appropriate reading; they used a delay of 200msec, by which time the contextually inappropriate probe word showed no residual priming.

While the fact that both meanings are initially primed is consistent with an autonomy position, this result is also completely consistent with an interactive account. Based on our earlier simulation with the ambiguous /d/-/t/ stimulus, this is just what we expect to see. Of course, the consistent slight advantage of contextually appropriate targets at 0 delay is also what we expect on an interactive-activation account. Further support for the idea that there is a context effect for 0-delay probes is provided by some observations of Simpson (1984), regarding another experiment by Onifer and Swinney (1981). He noted that Onifer and Swinney's experiments collected reaction times to probes for each meaning of an ambiguous word, both when the context favoured that meaning and when it favoured the alternative meaning. He then compared lexical decision times when the context was appropri-

ate, against lexical decision times when the context was inappropriate, and found that decision times were consistently faster with appropriate context.[3]

The fact that selection is complete at a longer delay is also fully consistent with the activation-competition processes that are assumed by the interactive activation approach; indeed the simulation shown in Fig. 1.2 is fully consistent with the pattern of results that we see in these experiments.

The initial advantage for contextually appropriate readings is small enough that it does not generally show up as significant. An interactive approach predicts that it should be possible to produce relatively strong contextual effects, even at short delays, when the context exerts relatively strong constraints. The question arises, then: Should we have expected the contexts used in these studies to produce strong effects? In general it is difficult to give a definitive answer to this question, since investigators have not tended to focus specifically on the degree of constraint.[4] The matter certainly deserves further scrutiny. However, there is one experiment that supports the prediction that relatively stronger contextual effects will be found early in processing when relatively strong contexts are used. An experiment by Simpson (1981) bears directly on this point. He selected a group of 60 ambiguous words and identified for each word a dominant and nondominant meaning. He then constructed five context sentences for each word, one that strongly favoured the dominant reading, one that weakly favoured the dominant reading, one that was neutral, one that weakly

---

[3]I should mention two somewhat countervailing caveats concerning the interpretation of data from these experiments. On the one hand, the response to the probe does not occur until several hundred milliseconds after the priming word, even when the probe follows the ambiguous word with 0 delay. Thus there is room for post-access processing of the ambiguous word before the response to the probe is made, even with a 0msec delay; an autonomy position could always take refuge in such a possibility to explain away effects of context at 0 delay. On the other hand, it has been noted that there may be some backward priming effects of the prime on the ambiguous word (Glucksberg, Kreuz, & Rho, 1986); this might have artificially raised the activation of the contextually inappropriate reading at 0 delay (but see Seidenberg et al., 1982).

[4]From an interactive activation point of view, predictability from the preceding context (i.e. *cloze* probability) provides a reasonable operational definition of degree of constraint; from the simulation with the input /dar?/, it was clear that even when there are only three possibilities consistent with the prior context, the context exerts primarily a selective, rather than a predictive effect. In this light, the predominantly selective pattern that is observed in the cross-modal experiments seems consistent with my own best guess about the predictiveness of the contexts used. In Swinney (1979), a single example stimulus is given in which there is a strongly constraining context. However, an examination of the full set of materials used by Onifer and Swinney (1981) indicates that in these later studies, at least, there was a wide range of contextual constraint. For example, consider the context: "The office walls were so thin they could hear the ..." It seems likely that subjects asked to guess would supply a variety of different continuations, with *ring*, the actual ambiguous word, being only one of many possibilities.
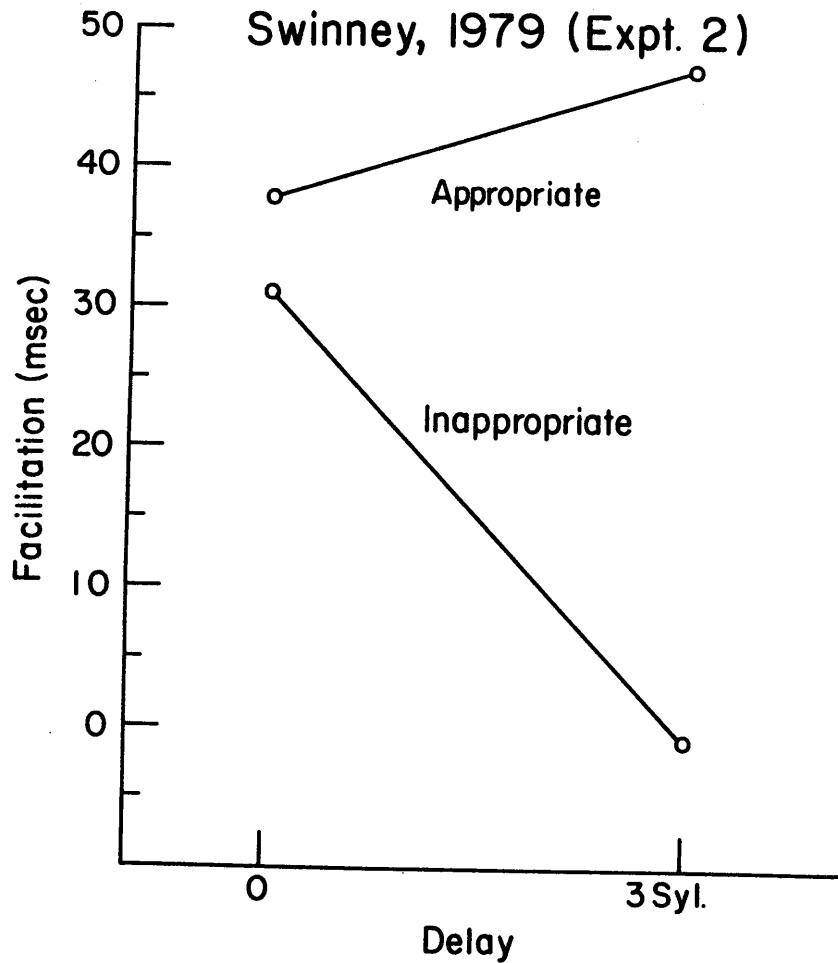
FIG. 1.4. Interaction of context and delay in the cross-modal priming experiment of Swinney, 1979.

TABLE 1.1
Priming Effects of Ambiguous Words in Context, 0 Delay

|	Appropriate Meaning	Inappropriate Meaning	A > I?
TLS 1979	33.5	22	YES
Swinney 1979			
Expt 1	70	50	YES
Expt 2	38	31	YES
STLB 1982			
Expt 3	17.5	13.5	YES
Expt 4			
(noun-verb)	16	28	NO
Expt 5	20	15	YES
MEAN	32.5	26.5	5 out of 6

FIG. 1.5. Effects of dominance and context from Simpson, 1981. Data from two groups of subjects are combined. One group received the strong and neutral contexts, and the other received the weak and neutral contexts. For the neutral conditions, I have connected the points through the mean averaged over the two groups. The horizontals at the top and bottom of the vertical bars represent the values obtained by the strong and weak context groups, respectively.

favoured the subordinate reading, and one that strongly favoured the subordinate reading. He presented these sentences to subjects, then followed the final word with a probe related either to the dominant or the subordinate meaning, or with a control, unrelated word. The probe occurred 120msec after the offset of the ambiguous prime word.

I have graphed the facilitation effects Simpson found in Fig. 1.5, as a

function of the strength of the context (from strongly favourable to the meaning related to the probe to strongly unfavourable) separately for the dominant and subordinate probes.

As the figure makes plain, there is a strong effect both of dominance and of context, as well as a context by dominance interaction. The interaction is such that when the context is strong, it completely wipes out the effect of dominance. Only when the context is weak or neutral is a strong dominance effect found.

The effects shown in this figure are exactly the kind of effects we would expect to find from an interactive activation model. Each of the two factors manipulated should produce an effect, but only when it is not dominated by the other factor. These kinds of effects are ubiquitous, as I have already noted, and are naturally accounted for by the principles of interactive activation. Unfortunately, there was a delay of 120msec after the ambiguous word in Simpson's experiment before the presentation of the probe; thus there is room to argue that the strong effects of context that he observed were due at least in part to this delay. Thus a definitive test of the predicted immediate context effect with strongly constraining contexts must await further research.

Thus far I have argued from characteristics of interactive activation mechanisms as observed in simulations of lexical effects on phoneme perception. Some readers may wonder whether these general characteristics of interactive activation mechanisms can actually be incorporated in a working model of meaning selection. In fact, both Cottrell (1985) and Kawamoto (Note 4) have developed simulation models that incorporate the principles of interactive activation and that exhibit effects in meaning selection that are analogous to those that I have described for the speech perception simulations. Kawamoto's model used distributed patterns of activation over an ensemble of units to represent the alternative readings of an ambiguous word, instead of the local representations that have been used in the interactive activation models of visual word perception and speech perception. In spite of this difference, his model produces the same kinds of effects that we have seen in other interactive activation models.[5]

I have argued that the results we have reviewed are consistent with the interactive approach, but I do not mean to suggest they cannot be accounted for within an autonomy position. One possible account for early context effects is to suggest that priming can occur within the lexical access mechanism itself. Indeed, Burgess, Seidenberg, and Tanenhaus (Note 1) accounted for the initial, selective access effects that were found in two of

---

[5]I would like to acknowledge here the contributions of Alan Kawamoto's work to this part of this article. His simulations and his review of the literature served as the basis for this discussion of lexical ambiguity resolution.

their experiments in terms of such effects. Intra-lexical priming might also be cited as a possible source of the advantage for contextually appropriate readings in other studies. Unfortunately, the case for this is far from clear at this point. No definitive studies have been done showing that contextual effects only result from intra-lexical factors, controlling for degree of constraint. It would seem that it behoves researchers on both sides of this debate to find ways of separating degree of constraint from intra- vs. inter-level source.

An autonomy account can also be salvaged if it is assumed that the observed priming effects reflect the results of post-access processes. Thus, as I stated at the outset, the finding that there are effects of context on responses to early probes is not compelling evidence against an autonomy account. My purpose has only been to show that the facts that have emerged from these cross-modal priming studies do not speak against an interactive position.

Let me note in closing that there are tests that can be done to test the interactive account. A strong test would be to examine whether context influences the activation of the meanings of an ambiguous word, even under conditions where it is strong enough to allow subjects to guess the identity of the ambiguous word quickly and correctly from the contextual information alone. In such a case interactive activation predicts that the inappropriate meaning will be less active at the earliest point that shows activation for either meaning.

Autonomy of Syntax

Recently, Lynne Frazier and her associates have proposed that syntactic processing is autonomous. In Frazier (Note 3), the suggestion is made that the syntactic processor initially makes decisions in terms of a very general principle known as minimal attachment, and provides a single parse to a "thematic processor" for acceptance or rejection. Here I am not so much concerned with the specific principle of minimal attachment per se, as with the more general claim that initial parsing decisions are unaffected by constraints arising from semantic/thematic considerations.[6] I will consider two experiments that have been taken as evidence for the autonomy position, both reported in Rayner, Carlson, and Frazier (1983). The first shows that plausibility based on knowledge of real-world constraints has little or no effect on the initial processing of reduced relative clauses attached to sentence

---

[6]I do not mean to take a particular stand on the exact characterisation of the higher-level factors that can be brought to bear on syntactic processing; by semantic-thematic constraints (henceforth, simply called *semantic*), I mean to include a range of constraints that arise from our knowledge of the meanings of words and of the ways the entities they refer to might plausibly be interrelated in the situations that we describe in sentences.

APPR-B*

intial noun phrases. The second shows a reading-time advantage for sentences containing a prepositional phrase that is minimally attached, compared to matched sentences in which the ultimate interpretation requires nonminimal attachment. I will discuss these in turn, dealing with the first one rather more briefly.

Reduced Relatives

In Rayner et al.'s (1983) first experiment, subjects read reduced relative sentences like the following:

The florist sent the flowers was very pleased.   1a

Such sentences, of course, have been well-studied since the early work of Bever (1970), who used them to support his argument for a particular sentence processing strategy he called the "NVN" strategy. According to the NVN strategy, a sequence that can be interpreted as noun-verb-noun, that is not otherwise marked as subordinate, is taken to specify an actor-action-object sequence. Phrases like "The florist sent the flowers" engage this strategy, and so lead to a garden-path effect, causing the subject to slow down and/or back up when information inconsistent with this effect is encountered.

That this NVN strategy is very potent in English is indicated by the fact that it is strong enough to completely over-ride semantic/thematic constraints. For example, adult English speakers asked to act out the sentence "The pencil kicked the cow" will pick up the pencil and knock over the cow with it, even though pencils are inanimate and therefore cannot ordinarily kick (Bates, McNew, MacWhinney, Devescovi, & Smith, 1982). Apparently, the NVN strategy is strong enough to over-ride semantic constraints in English.

It is important to my argument to note that, in other languages, syntactic constraints need not be so over-riding. For example, in Italian, there is a tendency to use the actor-action-object strategy in interpreting NVN sequences, but this tendency is not over-riding for Italians. Accordingly, Italians interpret analogues of "the pencil kicked the cow" in accordance with semantic constraints, even though they tend to treat the first noun as agent in more neutral sentences, such as "The horse kicked the cow" (Bates et al., 1982).

The point, so far, is that syntactic cues vary in strength from language to language, and there is no universal prepotency of syntax over semantics. It just so happens in English that there is a very strong tendency to treat NVN as actor-agent-object. In English, this particular syntactic cue is strong enough to over-ride semantic constraints such as animacy constraints on the agents of action verbs, as Bates et al. have shown.

In their Experiment 1, Rayner et al. (1983) compare reading times for reduced relative sentences like (1a) in which the NVN = actor-action-object reading of the beginning of the sentence seems very plausible with other sentences in which such a reading seems somewhat less plausible, such as (1b);

The performer sent the flowers was greatly pleased.   1b

Although performers can send flowers, they are less likely to do so than florists. Thus, one might reason, if subjects were able to make use of semantic constraints in on-line syntactic processing decisions, then they should not be as strongly misled in sentences like (1b). However, Rayner et al. found that subjects were slow to process the disambiguating portion of the sentences (in this case, "was greatly pleased"), regardless of the plausibility of the actor-action-object interpretation of the first NVN sequence, indicating that they were led down the garden path in both cases. Similar null effects of animacy of the sentence-initial noun-phrase or of preceding context have been reported by Ferreira and Clifton (1986).

Though the consistent lack of an effect in these cases might seem compelling at first sight, it is important to realise that it does not necessarily mean that syntactic processing decisions are unaffected by plausibility factors in all cases. We have reason to believe from other research that word order is very powerful as a cue in English, and that the NVN sequence is a compelling cue for an agent-action-object interpretation. In contrast, the plausibility manipulation used by Rayner et al. seems rather weak; for example there is no reason to suppose that a performer could not send flowers, say to a rival at the opening of a new show. My argument, quite simply, is that we cannot put weak cues against strong cues and expect that the weak cues will produce strong effects; indeed we have seen how strong cues can completely over-ride weaker ones in one of our initial illustrative simulations. We have independent evidence that demonstrates the potency of the NVN strategy, and so we cannot be surprised to find that weak contextual constraints have no reliable effects. The interactive activation framework makes clear that if we wish to find effects of a particular factor, we must look at situations in which there are no other factors exerting overpowering effects.

Prepositional Phrase Attachment

Just such a situation is provided by PP attachment ambiguities, such as the one that arises in sentences like "The boy hit the girl with the doll." In comprehending such sentences, the reader must decide whether to treat "the doll" as the instrument of hitting, thereby attaching it to the verb phrase; or whether to treat it as an object in the girl's possession, thereby attaching it a constituent of a complex noun-phrase headed by "the girl."

Such decisions are clearly influenced by thematic plausibility constraints. Consider, for example, the following sentences:

The spy saw the cop with binoculars.   2a
The spy saw the cop with a revolver.   2b

In the former sentence, we tend to treat "binoculars" as an instrument; in the latter, we treat "revolver" as a possession of the cop. In general, it appears that the verb and all of the noun phrases influence these decisions. Compare, for example,

The spy shot the cop with binoculars.   3a
The spy shot the cop with a revolver.   3b

and

The woodpecker saw the bird-watcher with binoculars.   4a
The bird-watcher saw the woodpecker with binoculars.   4b

Indeed, Oden (1978) has shown that attachment decisions can be influenced by the identities of the various NPs in the sentence and by preceding context.

No-one doubts the role of these constraints in the ultimate interpretations assigned to sentences. What is at issue is whether such constraints affect the initial attachment decisions subjects make in the course of reading or listening. An interactive account would assume that the initial attachment decision is susceptible to influence from semantic constraints: In view of the fact that both kinds of attachments are encountered frequently, there would be no reason to suppose that there would be a strong syntactic bias in favour of one attachment over the other. Frazier, however, has pointed out that the attachment of the preposition phrase as a constituent of the verb phrase would require the creation of no extra structure, and therefore she has proposed that verb-phrase (VP) attachment is tried first by the syntactic processor, independent of semantic constraints.

The second experiment reported by Rayner et al. (1983) addressed this claim. They presented subjects with sentences like (2a) and (2b), with an extra final clause added, and measured reading time as in their first experiment. They reasoned that, if the syntactic processor initially prefers VP attachments, then reading times should be slower for sentences like (2a), where a VP attachment turns out to be consistent with thematic considerations. The results of the experiment supported this prediction: Reading times were somewhat slower on and after the disambiguating word in the versions of the sentences where the ultimate reading favoured attachment of the prepositional phrase to the preceding noun-phrase (NP).

While the results were consistent with this prediction, it turns out that there is an alternative account. It is possible that the effects observed by Rayner et al. are not due to a syntactic preference for minimal attachment, but to the fact that, in Rayner et al.'s materials, there is a consistent semantic

bias in favour of the minimal completion. To show this, Taraban and McClelland (Note 5) asked subjects to read Rayner et al.'s sentences, through the preposition at the beginning of the critical prepositional phrase, and then to generate an expectation for the completion of this phrase. The subject then saw either the VP or the NP completion, and was asked to rate how well the actual completion matched the expectation. Subjects rated the VP completions significantly closer to their expectations, on average, than the NP completions (3.62 vs. 2.90 on a 5-point scale).

To determine whether it was this greater concordance with expectations that was determining the advantage for VP over NP completions, Taraban and McClelland constructed 20 additional sentence pairs that were intended to produce expectations favouring an NP completion. An example is:

I read the article in the ...   5a

This can be completed with a word like "magazine," in which case the PP is attached to the NP, or with a word like "bathtub," in which case the PP is interpreted by most subjects as being attached to the VP. The completion words used in the two conditions were matched over the set of materials for both length and frequency. As intended, the NP completions of Taraban and McClelland's sentences were rated closer to subjects' expectations than the VP completions (3.90 vs. 2.98).

Once ratings had been collected, both Rayner et al.'s sentences and Taraban and McClelland's new sentences were presented to another group of subjects in a word-by-word reading time task. At the beginning of each trial the subject pressed a button causing the presentation of a row of dashes, blanks, and punctuation marks. Each dash indicated the presence of a letter in the to-be-read sentence, with blanks indicating the spaces between words. The next press of the button caused the first set of blanks to be replaced with the first word of the sentence. Each subsequent press of the button caused the next word to be presented and the preceding word to be replaced with blanks. The last word of the sentence was always the disambiguating word. When the subject pressed the button after reading this word, a question appeared. Subjects were instructed to read the sentences as rapidly as possible consistent with good comprehension, and the answers to the questions were recorded by the experimenter. Accuracy was very high, and did not differ between experimental conditions. In addition to the 29 target sentences, there were 66 filler sentences. Seven of these were used to balance the frequency of NP and VP attachments of sentence final prepositional phrases. The remaining 59 were fillers of many different types included to vary the materials so that subjects would not get into a set of expecting a sentence-final prepositional phrase.

The reading times for the final words of the sentences are shown in Fig. 1.6a, broken down by attachment and source.

FIG. 1.6.   Opposite effects of attachment on reading time for target words triggering different attachment decisions, for sentences of Rayner et al. (1983) (RCF) and Taraban and McClelland (TM). In the first Experiment, (a), the sentence ended with the target word, and the reading times shown are for this word only. In the second experiment (b), the sentence continued on beyond the target word, and reading times are based on the sum of the time spent reading the target word and the three following words.

Two things are apparent from the results. First, with Rayner et al.'s materials, we were able to replicate their effect showing faster reading times for VP vs. NP attachments. Second, however, we found that with our materials, this effect was reversed, and reading times were actually shorter for NP completions than for VP attachments. There was no main effect of attachment type, but there was a highly reliable interaction of completion type with source (RCF vs. TM). There was also a main effect of source, but this is not interpretable, since Taraban and McClelland's completions were generally shorter and more frequent than those used by Rayner et al.

It has often been suggested that the time spent reading the final word of a sentence reflects extra, integrative processes that do not occur at other points. Thus, the reading times Taraban and McClelland observed in this experiment might reflect such integration effects, and these effects might be masking a real effect of attachment that would appear if it had not been overshadowed by such sentence-final integration effects. To address this problem, Taraban and McClelland extended the sentences. For the Rayner et al. sentences we used continuations they had used, and for our own we constructed completions of the same kind. In all cases, the continuation began with a conjunction that clearly indicated the beginning of a new clause, such as "while" or "because."

Figure 1.6b shows the total reading time for the target word and the following three words, broken down by VP vs. NP attachment and source. Once again there was no main effect of attachment, but there was a strong attachment by source interaction. Finally, Fig. 1.7 shows the difference in reading times between the VP and NP completions of the sentences, on a word-by word basis, starting with the disambiguating word.

The figure indicates that there is no effect of condition on the reading time
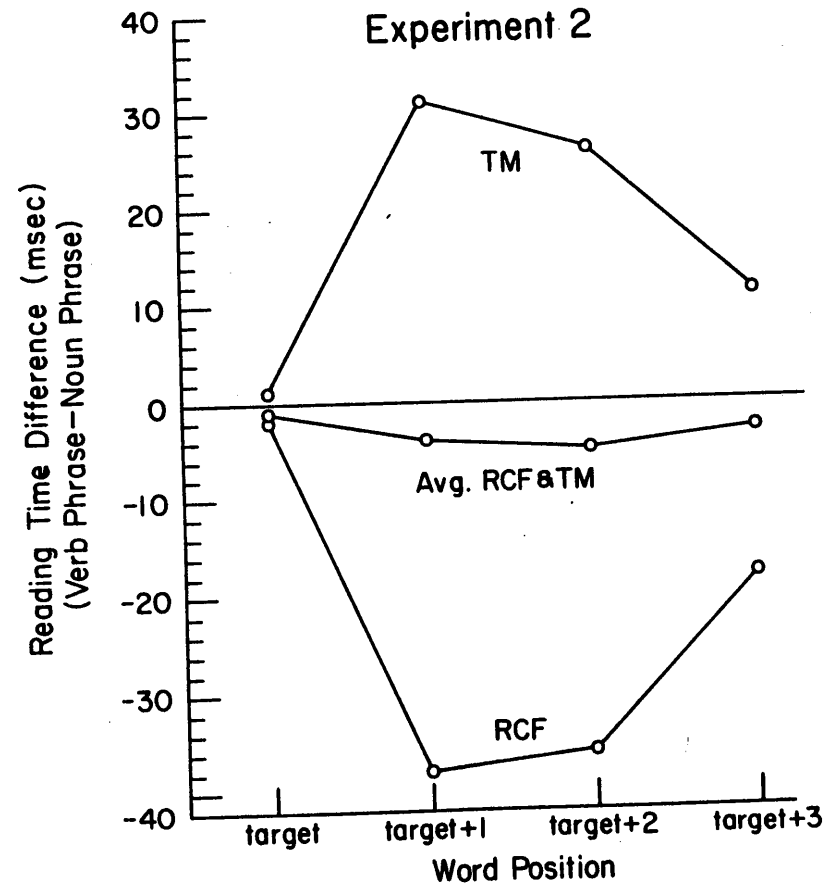

FIG. 1.7.   The time-course of the processing difference between NP and VP attachment versions of the Rayner et al. (RCF) and Taraban and McClelland (TM) sentences. Times shown are reading times for words in the NP-attachment version, minus reading times for words in the VP-attachment version, for the target word and each of the three following words.

for the disambiguating word itself. However, there is an effect in each of the next two words; by the third word after the disambiguation, the difference seems to have disappeared. It would appear from this analysis that processing that occurred on the disambiguating word when it was the last word of the sentence is being spread out over subsequent words in this case. As before, there is no evidence that this extra processing reflects a disruption that occurs with nonminimal completions in general. Rather, it appears that the extra processing occurs for minimal or nonminimal completions, depending on whether the VP or NP completion is closer to the subjects' expectations.

Once again, I do not intend to suggest that the facts actually rule out the autonomous syntax position in favour of an interactive view; it remains possible to suppose that syntactic processing is autonomous, but that what is determining the reading times we are observing is not (or not simply) the output of this syntactic process. On the other hand, the interactive activation approach deserves some credit for giving us guidance in the search for cases in which processing times appear to be dominated by semantic as opposed to syntactic considerations. At the very least it seems clear that Rayner et al.'s second experiment provides little reason to doubt that semantic considerations can play a role in syntactic decisions, given the fact that it appears to be semantic and not syntactic factors that are controlling reading times for these sentences.[7]

In summary, I would suggest that the findings of Rayner et al. need not be interpreted as favouring any version of autonomous syntax hypothesis. Though syntactic cues are sometimes so strong that they overshadow semantic constraints, we find that under other conditions semantic constraints do appear to exert relatively immediate effects.

DISTINGUISHING INTERACTIVE FROM AUTONOMOUS PROCESSING

Although some quibbling may be possible, the evidence appears to me to be fairly clear in supporting the following proposition: Decisions about representational units of all kinds involve the consideration of multiple sources of information.

---

[7]The fact that we used a word-by-word reading time measure, coupled with the fact that our effects only show up on the word *after* the disambiguating word, might be taken as evidence that in fact the effects we observed occur *after* an initial syntactic attachment process that works immediately and is reflected only in eye fixation duration. In this context it should be noted that Rayner et al.'s findings did not show up clearly in fixations on the target word; indeed the statistical evidence for their effect was somewhat weak in their eye-movement data, perhaps because subjects tend to overlap the completion of higher levels of processing with the intake of subsequent words.

However, this can be seen simply as a restatement of some of the basic findings, rather than as a statement about whether the processing system is inherently interactive or not. To see this, I will briefly consider two cases: the lexical effect on phoneme identification (Ganong, 1980) and the role of semantic context in resolving the attachment ambiguities we have been discussing. In both cases, we might account for the results with a purely bottom-up processing system, in which each module operates completely independently of influences from higher levels of processing. Thus in Ganong's case, one may propose that the phoneme level passes to the word level activations indicating which phonemes are consistent with the input and to what extent; and that the word level uses these graded activations, in conjunction with lexical constraints, to determine which word(s) are consistent with the input. Thus if a phoneme ambiguous between /g/ and /k/ is heard, the phoneme level may pass on the ambiguity to the word level. Ganong's finding could simply result from choosing as an overt response the phoneme that is most consistent with the word that the subject has heard. The decision is still based on information from multiple sources, but this integration of information does not occur at the phoneme level of processing within the perceptual system; instead, it occurs in some later decision-making process that can consult the final output of the word level.

In the sentence processing case, the situation is analogous. One could suppose that the syntactic processing mechanisms operate autonomously, passing on to higher levels the output of a preliminary syntactic analysis. In the case of attachment ambiguities such as those considered here, one might assume (contrary to Frazier, but more or less consistent with the recent view of Marcus, Hindle, & Fleck, 1983) that the output reflects the possible attachments that are consistent with the syntax, with each activated to a degree that reflects its relative likelihood based on syntactic considerations. The semantic processor could then make use of this information, in conjunction with semantic constraints, to achieve an interpretation that was jointly constrained by syntactic and semantic factors.

This purely bottom-up story has many of the same implications as an interactive account, since it explains how influences from all levels can have effects on the final outcome of processing. It is certainly consistent with a large number of existing experiments on contextual influences. One might ask, then, whether there is any way of distinguishing this purely bottom-up account from an interactive view.

Fodor (1983) has made one suggestion. He has observed that to counter unidirectional accounts, it is necessary to show "that the information fed back interacts with interlevels of input-processing and not merely the final results of such processing." Thus, for example, if one could show that the results of semantic processing are fed back into the syntactic processor in such a way as to influence subsequent syntactic processing decisions, or that

the results of lexical processing are fed back into the phonetic level so as to influence subsequent phonetic processing decisions, then one would have provided evidence that processing is indeed interactive.

To illustrate this approach, I will describe a recent experiment by Elman and McClelland (Note 2). In this experiment, we relied upon the fact that listeners compensate for coarticulatory influences of one speech sound on the acoustic realisation of neighbouring sounds. In the case we exploited, the phonemes /s/ and /S/[8] alter the acoustic realisation of a subsequent /t/ or /k/; listeners compensate for this coarticulation effect by adjusting the perceptual boundary between /t/ and /k/, so that a sound that would be on the boundary in a neutral context tends to be heard as a /k/ when it occurs just after a /s/, but as a /t/ when it occurs after a /S/. We reasoned as follows. First, we assumed that this coarticulatory compensation is an intrinsic characteristic of processing at the phoneme level. Given this, we noted that it should be possible to use lexical constraints to get subjects to interpret a sound halfway between /s/ and /S/ as a /s/ in one context and as a /S/ in another. Now if, as we assumed, this lexical effect operates by feeding back activation to the phoneme level; and if, as we also assumed, interactions at the phoneme level are responsible for the coarticulatory compensation effect, then the lexical effect on the ambiguous /s/-/S/ sound should trigger a coarticulatory compensation effect that influences the phonetic interpretation of an ambiguous /k/-/t/ sound. On the other hand, if Ganong's effect operates only on the final results of phonetic processing, and does not feed back anything to the phonetic level, then we would expect no coarticulatory compensation as a result of the lexical effect.

We therefore took pairs of words (e.g. "tapes/capes") distinguished by initial /t/ vs. /k/ (or /d/ vs. /g/, which exhibit the same effects of preceding /s/ and /S/) and constructed from recorded tokens of these words a set of seven stimuli beginning with sounds varying between /t/ and /k/ in small steps. Each of these stimuli was preceded by one of two context words. In one experiment, one word (e.g. "foolish") actually ended in /S/ and the other (e.g. "Christmas") actually ended in /s/. In another experiment, the same context words were used but the final segments were replaced by an ambiguous sound that was determined in pre-testing to fall halfway between /s/ and /S/, here designated as /?/.

The first experiment simply replicated the coarticulatory influence of /s/ and /S/ on the identification of borderline /t/-/k/ stimuli, as previously described by Mann and Repp (1982); as expected, words ending in /s/ tended to lead to an increased probability of /k/ responses to the subsequent /t/-/k/ stimulus, while the words ending in /S/ tended to lead to an increased probability of /t/ responses.

---

[8]I use /S/ to stand for the "sh" sound in "ship."

The second experiment provided the crucial test for the interaction hypothesis. Here, we found that prior context did indeed trigger coarticulatory compensation for the lexically-determined /s/ or /S/ phoneme; for example, subjects reported /k/ more often after "Christma?" than after "fooli?", just as predicted. The results for several context/target sets involving /t/-/k/ and /d/-/g/ identification are shown in Fig. 1.8.
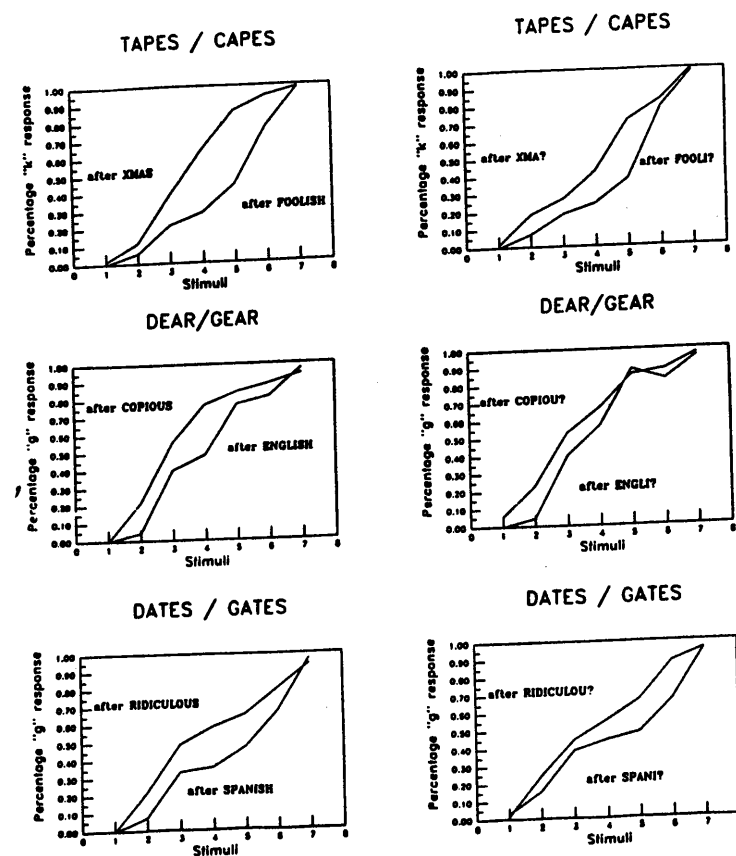

FIG. 1.8.   Identification curves for three sets of experimental stimuli used by Elman and McClelland (1986). The left panels show the effects of acoustically distinct "s" and "sh" sounds on /t/-/k/ and /d/-/g/ judgements; the right panels show the effects of acoustically identical (lexically disambiguated) sounds halfway between "s" and "sh" (represented by ?). The label above each panel indicates the words that were used to bracket the ambiguous /t/-/k/ and /d/-/g/ stimuli; the labels associated with each curve indicate the preceding context for the judgement percentages (percentage /g/ or /k/ judgements, depending on the continuum) indicated by the corresponding curve.

The results of this experiment demonstrate that lexical influences on phoneme identification can induce coarticulatory compensation, as predicted from the interaction hypothesis. This is exactly what we would expect if, indeed, feedback from the lexical level actually does influence processing at the phoneme level, rather than simply influencing the interpretation of the outcome of such processing. More importantly, the experiment demonstrates a method that I think holds some considerable promise of providing a way of determining the extent of interaction in perceptual and linguistic processing.

It remains possible to salvage a bottom-up account for these findings, but I do not think this is a very attractive option. To do so, one must suppose that compensation for coarticulation is accomplished by the same "late" mechanism that uses lexical information to make decisions about the identity of phonemes. This seems an unattractive suggestion, because compensation for coarticulation is so often taken as an intrinsic and basic function of the mechanisms of phoneme perception (see, for example, Liberman, Cooper, Shankweiler, & Studdert-Kennedy, 1967). To ascribe this function to some "later" level would be to deprive the machinery of phoneme perception of one of its most crucial roles; or to duplicate needlessly the intricate knowledge of coarticulatory influences that is assumed to be present in the mechanisms of phoneme perception in mechanisms of post-perceptual judgement.

More generally, it would always be possible to say that processing interactions that are assumed to result from intra-level influences were actually occurring at a higher level, and thereby to sidestep any possible applications of Fodor's suggested test. But this step is only palatable, it seems to me, if the higher-level decision can be made using information that would ordinarily be assumed to be available to the higher level. This, it seems quite sensible to suppose that phonetic ambiguity could be passed up to a later stage for resolution at the word level provided the word level does it by using lexical constraints. But if the word must use the very sorts of information usually attributed to the phoneme level, then the entire notion of encapsulation of knowledge is undermined.

This discussion brings up another point, and that is, why bother with feedback? What's the good of it? Why should it matter if higher levels feed back information into lower levels? Why should they not simply resolve the ambiguities that are passed on to them whenever they can, and forget about providing feedback supporting one alternative over the other?

The good of feedback is that it permits processing on lower levels to be guided from above, thereby allowing them to provide higher levels with better information. Our coarticulation study gives one example of this. If higher levels can help lower levels decide on the identity of phonemes that are perceptually indistinct, then lower levels can use this information to adjust for coarticulation better than they could otherwise. Similarly, at the syntactic level, if higher levels can influence the formation of syntactic representations of one constituent, they will allow the syntactic level to be better prepared to provide the best analysis of what will come later on in the sentence. In both cases, this allows the lower level to do a better job in providing information to the higher level.

SUMMARY

In the preceding sections of this paper, I have described a framework for modeling the process of forming representations in processing written and spoken language. I have shown how this framework can help us understand why contextual effects may be obtained under some circumstances and not others, and why it often appears to exert selective, as opposed to predictive, effects.

In the course of making these observations, I have argued that some of the evidence that has been taken in support of the idea that lexical access and syntactic processing are invulnerable to external influences is fully consistent with an interactive account. I do not say that this part of the analysis proves that the autonomy position is wrong, only that several of the reasons that have been given for believing that it is wrong are far from compelling.

Finally, I have indicated that there is hope of finding empirical evidence relevant to distinguishing between interactive and feed-forward accounts of information processing: Such evidence takes the form of demonstrations that higher levels of processing can trigger processes at lower levels, increasing the quality of the results they pass on later to higher levels.

It remains to build explicit models of interactive processing at higher levels. Of course, this is a difficult task for any processing framework; certainly no adequate model of the formation of a representation of the event or scene described by a sentence has been proposed to date. From what we know about the susceptibility of higher levels of language processing to contextual information (cf. Bransford & Johnson, 1973), it seems fairly clear to me that any adequate model will have to incorporate the principles of interactive activation. What is not clear at this point is how these principles will need to be elaborated and supplemented to capture the structural complexities that arise at higher levels. This remains a central issue for future research.

REFERENCES

Bates, E., McNew, S., MacWhinney, B., Devescovi, A., & Smith, S. (1982). Functional constraints on sentence processing: A cross-linguistic study. *Cognition, 11*, 245–299.

Bever, T. G. (1970). The cognitive basis for linguistic structures. In J. R. Hayes (Ed.), *Cognition and the development of language.* New York: Wiley.

Bransford, J. D. & Johnson, M. K. (1973). Considerations of some problems of comprehension. In W. G. Chase (Ed.), *Visual information processing.* New York: Academic Press, 383–438.

Cottrell, G. (1985). *A connectionist approach to word sense disambiguation* (TR-154). Rochester, New York: University of Rochester, Department of Computer Science.

Elman, J. L. & McClelland, J. L. (1986). Exploiting the lawful variability in the speech wave. In J. S. Perkell & D. H. Klatt (Eds.), *Invariance and variability of speech processes.* Hillsdale, N.J.: Lawrence Erlbaum Associates Inc.

Ferreira, F. & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language, 25,* 348–368.

Fodor, J. A. (1983). *Modularity of mind: An essay on faculty psychology.* Cambridge, Mass.: M.I.T. Press.

Ganong, W. F. (1980). Phonetic categorisation in auditory word perception. *Journal of Experimental Psychology: Human Perception and Performance, 26,* 110–115.

Glucksberg, S., Kreuz, R. J. & Rho, S. H. (1986). Context can constrain lexical access: Implications for models of language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 323–335.

Grossberg, S. (1976). Adaptive pattern classification and universal recoding: Part 1. Parallel development and coding of neural feature detectors. *Biological Cybernetics, 23,* 121–134.

Grossberg, S. (1978). A theory of visual coding, memory, and development. In E. L. J. Leeuwenbert & H. F. J. M. Buffart (Eds.), *Formal theories of visual perception.* New York: Wiley.

Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart, J. L. McClelland, & the PDP research group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. I.* Cambridge, Mass.: Bradford Books.

Huey, E. B. (1968). *The psychology and pedagogy of reading.* Cambridge, Mass.: M.I.T. Press. (Reprinted from Macmillian Company, 1908.)

Liberman, A. M., Cooper, F. S., Shankweiler, D., & Studdert-Kennedy, M. (1967). Perception of the speech code. *Psychological Review, 84,* 452–471.

Mann, V. A. & Repp, B. H. (1982). Fricative-stop coarticulation: Acoustic and perceptual evidence. *Journal of the Acoustical Society of America, 71,* 1562–1567.

Marcus, M. P., Hindle, D., & Fleck, M. M. (1983). D-theory: Talking about talking about trees. *Proceedings of the Association for Computational Linguistics.*

Marslen-Wilson, W. D. & Welsh, A. (1978). Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology, 10,* 29–63.

McClelland, J. L. (1985). Putting knowledge in its place: A scheme for programming parallel processing structures on the fly. *Cognitive Science, 9,* 113–146.

McClelland, J. L. (1986). The programmable blackboard model of reading. In J. L. McClelland, D. E. Rumelhart, & the PDP research group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. II.* Cambridge, Mass.: Bradford Books.

McClelland, J. L. & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology, 18,* 1–86.

McClelland, J. L. & Kawamoto, A. H. (1986). Mechanisms of sentence processing: Assigning roles to constituents. In J. L. McClelland, D. E. Rumelhart, & the PDP research group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. II.* Cambridge, Mass.: Bradford Books.

McClelland, J. L. & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 88,* 375–407.

Mozer, M. (1987). Early parallel processing in reading: A connectionist approach. In M. Coltheart (Ed.), *Attention and performance XII.* London: Lawrence Erlbaum Associates Ltd.

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton-Century-Crofts.

Oden, G. C. (1978). Semantic constraints and judged preference for interpretations of ambiguous sentences. *Memory and Cognition, 6,* 26–37.

Oden, G. C. & Massaro, D. W. (1978). Integration of featural information in speech perception. *Psychological Review, 85,* 172–191.

Onifer, W. & Swinney, D. A. (1981). Accessing lexical ambiguities during sentence comprehension: Effects of frequency on meaning and contextual bias. *Memory and Cognition, 9,* 225–236.

Rayner, K., Carlson, M., & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *Journal of Verbal Learning and Verbal Behaviour, 22,* 358–374.

Rumelhart, D. E. (1977). Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and performance VI.* Hillsdale, N.J.: Lawrence Erlbaum Associates Inc.

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In D. E. Rumelhart, J. L. McClelland, & the PDP research group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1.* Cambridge, Mass.: Bradford Books.

Rumelhart, D. E. & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception: Part 2. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review, 89,* 60–94.

Samuel, A. G. (1981). Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General, 110,* 474–494.

Seidenberg, M. S., Tanenhaus, M. J., Leiman, J. M., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology, 14,* 538–559.

Simpson, G. B. (1981). Meaning dominance and semantic context in the processing of lexical ambiguity. *Journal of Verbal Learning and Verbal Behaviour, 20,* 120–136.

Simpson, G. B. (1984). Lexical ambiguity and its role in models of word recognition. *Psychological Bulletin, 96,* 316–340.

Smolensky, P. (1986). Neural and conceptual interpretation of PDP models. In J. L. McClelland, D. E. Rumelhart, & the PDP research group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. II.* Cambridge, Mass.: Bradford Books.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behaviour, 18,* 645–659.

Tanenhaus, M. K., Leiman, J. M., & Seidenberg, M. S. (1979). Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behaviour, 18,* 427–440.

Waltz, D. L. & Pollack, J. B. (1985). Massively parallel parsing. *Cognitive Science, 9,* 51–74.

REFERENCE NOTES

1. Burgess, C., Seidenberg, M., & Tanenhaus, M. K. (1986). *Nonword interference and lexical ambiguity resolution.* Paper presented at the Program for the 27th annual meeting of the Psychonomic Society, New Orleans, LA.

2. Elman, J. L. & McClelland, J. L. (1987). *Cognitive penetration of the mechanisms of perception. Compensation for coarticulation of perceptually restored phonemes.* Manuscript submitted for publication.

3. Frazier, L. (1986). *Theories of sentence processing.* Manuscript.

4. Kawamoto, A. H. (1985). *Dynamic processes in the (re)solution of lexical ambiguity.* Unpublished doctoral dissertation, Brown University.
5. Taraban, R. & McClelland, J. L. *The role of semantic constraints in interpreting prepositional phrases.* Manuscript in preparation.