

Magnuson, J. S. (2008). Nondeterminism, pleiotropy, and single word reading: Theoretical and practical concerns. In E. Grigorenko & A. Naples (Eds.), *Single Word Reading*, pp. 377-404. Mahwah, NJ: Erlbaum.

Reader beware: OCR-based PDF with some errors.

CHAPTER EIGHTEEN

Nondeterminism, Pleiotropy, and Single-Word Reading: Theoretical and Practical Concerns

James S. Magnuson

University of Connecticut and Haskins Laboratories

It's daunting to be asked to write a chapter based on your lack of expertise. My primary expertise is in SWR – but *spoken word recognition* rather than *single word reading*. I was asked to read the chapters collected in this book, and comment from the perspective of a distinct, though closely related, field. Within this volume, rather stunning breadth and depth are represented. Consider the list:

Two chapters focus on developmental theories of literacy (Seymour) and spelling (Pollo, Treiman, & Kessler). Five chapters – two on orthographic-phonological consistency (Grainger & Ziegler; Kessler, Treiman, & Mullennix), and one each on developmental interactions between phonological and orthographic representations (Goswami), the role of semantics at the single word level (Keenan & Betjemann), and a new index of semantic richness that appears to moderate morphological effects (Feldman & Basnight-Brown) – provide landmarks for how deep and complex the questions have become in behavioral studies of single word reading. One chapter focuses on a powerful new paradigm of using artificial language materials to provide manipulations that would be virtually impossible with natural materials, and simultaneously provide a glimpse into acquisition processes (Hart & Perfetti; see also Mauer & McCandliss). Two focus on identification and remediation of dyslexia and reading disability (Piasta & Wagner; Royer & Walles). Four chapters review the remarkable progress over the past few years in mapping the brain regions and circuits that are crucial for reading (Frost, Sandak, Mencl, Landi, Moore, Della Porta,

Rueckl, Katz & Pugh; Mauer & McCandliss; Nazir & Husckauf; Simos, Billingsley-Marshall, Sarkari, & Papanicolaou). Finally, three chapters (Barr & Couto; Grigorenko; Olson) review how tantalizingly close the field is to identifying genetic and environmental influences on reading development (though Grigorenko provides some cautions about interpreting this work).

I was struck by two general themes that apply to the enterprise of studying single word reading. The first is *pleiotropy*, a theme of Grigorenko's chapter. Grigorenko introduces pleiotropy as follows: "...from the Greek *pleio*, or many, and *tropos*, manner, assumes 'multiple impact,' indicating diverse properties of a single agent, or that a single cause might have multiple outcomes." While this term is used technically to refer to effects of a single gene on multiple, possibly unrelated phenotypic characteristics, it resonates with similar notions at different levels of analysis from developmental psychology (equifinality), dynamical systems theories (coupling and emergentism), and more generally of the rampant nondeterminism, or many-to-many mappings, apparent in the information processing that underlies language understanding. The second issue is much more concrete: how to grapple empirically and theoretically with the multiple, interacting assumptions in theories of language processing.

So my chapter has two parts. In the first, I will argue that significant caution is warranted in focusing on any theorized level or stage of language processing (such as single word reading), on two bases: first, there is substantial evidence for radical interaction in language processing, and second, principles of theoretical neuroscience support the need for interaction in both processing and learning for a nondeterministic domain such as language processing. In the second part, I will argue that artificial language materials and, more importantly, computational models provide means for addressing the "material dilemma" of psycholinguistics – the ever-growing number of factors that must be controlled, which makes experimental design a truly Sisyphean task.

PART 1: PLEIOTROPY AND NONDETERMINISM

An obvious point of critique of the topic of single word reading is the focus on single words. Several chapters in this volume begin with justifications for focusing on single words. Hart and Perfetti argue that single word decoding is perhaps as close as one can come to a skill that is purely associated with reading. In contrast, there are other general cognitive traits (e.g., working memory, attention and motivation) that are argued to moderate rather than mediate broader aspects of reading, such as text comprehension. Maurer and McCandliss argue that focusing on single words makes pragmatic sense, as single word reading represents a "critical component process within reading" that can also provide a window into lower-level component processes with careful experimental designs. Olson points out that deficits in single word reading are highly correlated with deficits in reading in general, and Piasta and Wagner provide considerable support for the notion that the single word is the "nexus of ... reading problem(s)" (this nexus is held to be complex, as it is the junction at which influences of auditory processing, phonology, morphology, etc., can be simultaneously observed). These "proof is in the pudding" arguments

are compelling. There is no doubt that studies of single word reading have shed tremendous light on reading (indeed, studies of single word reading are arguably the source of nearly all our crucial knowledge of reading mechanisms).

In preparing to critique this focus on single words, I am reminded of a possibly apocryphal story about a prominent cognitive neuroscientist. After a conference talk about single word reading, he was asked what the impoverished, artificial tasks used in word recognition studies really have to do with reading outside the laboratory. The reply: "Nothing, unless one supposes reading involves words."

All the same, I will pursue this line of criticism. Let me reiterate that I will not challenge the usefulness of the single word level. While one might review the dangers of focusing on single words (e.g., the need for assessments of more complex reading tasks to ensure word-level findings generalize [cf. Royer & Walles, this volume, who report that word recognition can be boosted independently or even to the detriment of broader reading ability], or the fact that reading in context is demonstrably different from reading isolated words, at least for poor readers; e.g., Nicholson, 1991; see also Landi, Perfetti, Bolger, Dunlap, & Foorman, 2006), these issues are covered several times in preceding chapters. Rather, I will address two theoretical concerns for single word research. The first is that whether one considers the single word (or any other level of analysis) to be an isolable unit (a) based on theoretical principles identifying words as a discrete level of representation *or* (b) simply because it is useful pragmatically in the practice of research, the focus on a single level *functions* as a theoretical assumption and influences the types of research questions one asks. The second concern is that it may actually be the case that we can appeal to theoretical principles to assert that the single word level is discrete, or at least, not immediately subject to top-down interaction, and so can be treated as a distinct stage. I will argue that this position is untenable, based on theoretical principles that demand interaction in order to learn hierarchical representations like those involved in language processing, and based on strong evidence for early interaction in processing and learning.

First Concern: Implications of the Division of Labor in Psycholinguistics

The second principle is that of division into species according to the natural formation, where the joint is, not breaking any part as a bad carver might. –
(Plato's *Phaedrus*, 265e)

The justifications for focusing on single words offered in several chapters suggest a degree of unease with the focus on single words. There is a bit of tension on this point apparent in the chapter by Feldman and Basnight-Brown describing the impact of semantic richness in morphological characteristics, and in the arguments presented by Keenan and Betjemann against neglecting semantics and comprehension even at the single word level. Is the single word akin to a natural kind, neatly jointed and easily segmented from other aspects of language and cognition? Or is the focus a benign simplifying assumption that allows significant progress? If so, might it do so at the cost of misleading us somewhat about the larger reading and language processing systems?

Consider Marr's (1982) notion of a computational information processing theory. At the computational level of analysis, the focus is the computations performed by the system in a broad sense: what is the basic nature of the input and the output, and what general constraints can be identified in the mapping between them? Now consider a broad psycholinguistic domain, such as spoken or written language understanding. A computational theory of such a broad domain is intractable in a fairly transparent way – how do we tackle directly the problem of mapping from orthography or acoustics to message? Unsurprisingly, in the practice of psycholinguistics, we break the problem into several more tractable components. Psycholinguists working on speech perception have the job of determining how listeners achieve the mapping from the speech signal to consonants and vowels. Psycholinguists working at the level of spoken word recognition have the job of figuring out how listeners parse the phoneme stream into sound forms that provide access to the lexicon. Other psycholinguists have the jobs of figuring out how streams of word forms could be parsed into syntactic, semantic, and discourse structures.

For some researchers, these divisions correspond closely to theoretically motivated divisions of labor within the actual language processing system, in which hierarchical stages of processing provide fast, veridical perception and understanding (Norris, McQueen, & Cutler, 2000; Frazier & Fodor, 1978).

For others, the divisions are accepted as benign simplifying assumptions that have afforded progress at higher levels before all lower-level problems are solved. Proponents of highly interactive theories may find dubious the notion of discrete levels of representation and processing that correspond closely to linguistic levels of description. Rather, they might consider, for example, the assumption of phonemic input to word recognition, and even the notion of a specific mechanism that could be labeled word recognition, as useful but temporary and heuristic solutions. I will argue that it does not matter whether one adopts division of labor assumptions of convenience or principle; both reify the division of labor. That is, in spite of their obvious usefulness, division of labor assumptions are not benign fictions, but function like theoretical assumptions, constraining the sorts of research questions that are asked. To illustrate this point, consider the embedded word problem in spoken word recognition.

The division of labor assumption that we can consider a phoneme string to approximate the output of speech perception allows us to defer the perennially unsolved problems of speech perception, such as *lack of invariance* (the many-to-many mapping of acoustics to perceptual categories) and the *segmentation problem* (the lack of discrete boundaries between coarticulated phonemes). This allows research on spoken word recognition to focus on word-level problems, such as the embedding problem. McQueen, Cutler, Briscoe and Norris (1995) estimated that 84% of words in English have one or more words embedded within them (e.g., depending on dialect, *catalog* contains *cat*, *at*, *a*, *cattle*, *law* and *log*). This presents a significant theoretical challenge: how is it that all those embedded words are not recognized when we hear *catalog*? However, several recent results suggest the embedding problem is overestimated by division of labor simplifications.

While a few studies have shown that lexical access is exquisitely sensitive to fine-grained, subphonemic detail in the speech signal (Andruski, Blumstein, & Burton, 1994; Dahan, Magnuson, Tanenhaus, & Hogan, 2001; Marslen-Wilson & Warren, 1994), their focus typically has more to do with questions of how lexical competition is resolved. The larger implications were missed until quite recently. Davis, Marslen-Wilson, and Gaskell (2002) and Salverda, Dahan, and McQueen (2003) noted that vowel duration is inversely proportional to word length, such that the /ae/ in *ham* is, on average, slightly longer than that in *hamster* (by about 20 msec at typical speaking rates; cf. Lehiste, 1972; Peterson & Lehiste, 1960; Port, 1981). Both groups found that listeners are sensitive to these subphonemic differences (Davis et al. found evidence via priming, while Salverda et al. found more direct evidence using eye tracking). This suggests the embedding problem is not so dire: goodness of fit to lexical representations is graded according to fine-grained bottom-up details of the speech signal. *Ham* and *hamster* compete, but much less than *hammer* and *hamster*, which have more similar vowels in their first syllables.

Converging evidence has been found in visual word recognition, where there is an astounding sensitivity to fine-grained *phonetic* tendencies. For example, the longer it takes to pronounce a word (e.g., the pronunciation of *plead* tends to be longer than that of *pleat* due to durational effects of voicing; Peterson & Lehiste, 1960; Port, 1981), the longer it takes to process it even in a silent reading task (Abramson & Goldginer, 1997; Lukatela, Eaton, Sabadini, & Turvey, 2004). Lukatela et al. argue that this proves not only that visual lexical access is initially phonological (Lukatela & Turvey, 1994a, 1994b), but that lexical access is organized according to reliable *phonetic* patterns.

It is noteworthy that subcategorical sensitivity had been exploited previously in service of, e.g., asking questions about lexical competition (Dahan et al., 2001; Marslen-Wilson & Warren, 1994), and the vowel duration differences described decades earlier were well-known, but the joint implications of these findings for issues like the embedding problem were not immediately apparent. This is because the simplifying assumption that *the relevant grain for thinking about the input to spoken word recognition is the phoneme* has a powerful influence on the sorts of questions one asks and the patterns of results one is prepared to discover; that is, functionally, it acts as a computational-level theoretical assumption. Specifically, division of labor assumptions reify the stage view of processing, where each step in a series of processes creates a discrete product to be passed on to the next level.

This notion of *product* calls for *bidirectional* caution. It can lead us not only to underestimate the detail of the bottom-up input still available at some mid-level of description, but also to neglect the role of top-down interaction. That is, by assuming modular organization, even temporarily, we can ironically overestimate the complexity of processing required at any hypothesized level because we neglect the possibly helpful role of top-down constraints. However, there are long-running debates in psycholinguistics between proponents of autonomous processing theories (staged processing with feedforward encapsulation, such that low-levels are protected from feedback; Frazier & Fodor, 1978; Frazier & Rayner, 1982; Norris et al., 2000) and proponents of interactive models (e.g., MacDonald, Pearlmuter, &

Seidenberg, 1994; McClelland, Mirman, & Holt, in press), as well as debates about modularity more generally. There are also well-known studies that appear to support autonomous architectures. The goal of the next section is to analyze the case against and for interaction.

Second Concern: The Need for Interaction

The preceding argument has little implication for single word reading unless there are bidirectional constraints on the single word level – that is, that there is top-down interaction. If there is not top-down interaction, context dependence at the word level would be manageable, and there would be little concern about misconstruing the problem of language understanding by focusing on the single word level. So: how can we evaluate whether there is top-down interaction?

Consider again the concepts of pleiotropy, and more generally, equifinality. When one measures the state of any component of a complex system (i.e., any system of coupled, multiple parts that interact nonlinearly) at any time scale, one cannot recover the history of the isolated component – that is, the previous set of states of this component – without referring to states of the larger system. Multiple previous states can result in a single state (equifinality), but the converse is generally true as well: multiple outcomes can follow the same initial state, at least when what is known about the initial state is limited to only a portion of the entire system. This is true at multiple scales; genetically, developmentally, and perceptually and cognitively. A familiar perceptual/cognitive example is the context dependence of the H/A ambiguity shown at the left in Figure 18.1. Perception of the figure cannot be understood at the single letter level, but instead depends on lexical context. A slightly more complex example is shown at the right in Figure 18.1 (modeled after an example used by Friston, 2003). If the sentences are presented in isolation, readers experience little ambiguity in the “Jack and Jill...” example, despite the fact that “went” has been replaced with “event.” For this example, lexical and letter ambiguities are resolved at the sentence level.

These examples, by themselves, do not demonstrate on-line perceptual interaction, only that top-down and bottom-up information is integrated eventually. Whether the temporal locus of integration is “early” (perceptual) or “late” (post-perceptual) has been the subject of vigorous debate in several areas of psycholinguistics, and psychology and cognitive neuroscience more generally (e.g., Fodor, 1983), and many readers will be familiar with the logical arguments and empirical results that are typically marshaled against interaction. I will first review arguments for autonomous architectures and a pair of well-known results that are often cited in support of autonomous theories. I will provide an alternative account of the evidence, and then turn to principles of theoretical neuroscience to make a case that unsupervised learning of linguistic systems requires feedback.

Arguments Against Interaction. There are three primary arguments made in favor of autonomous architectures. First, *interaction is inefficient*. One of the clearest versions of this argument was made by Norris et al. (2000), who argued that

THE *Jack and Jill went up the hill.*
 CAT *The marathon was the last went*

FIGURE 18.1 A well-known letter ambiguity resolved at the word level (left), and a letter *and* lexical ambiguity resolved at the sentence level (right; based on an example used by Friston, 2003).

lexical-phonemic feedback could not possibly improve spoken word recognition. The logic is similar to specific principles of information theory (e.g., *law of diminishing information*; Kahre, 2002) and control theory (e.g., Ashby's [1962] *law of requisite variety*), which can be paraphrased as follows: once a signal enters a processing system, the information in the signal itself cannot be increased; at best it can remain constant. Norris et al. argue that this logic implies that a processing system with the goal of mapping, e.g., from phonemes to words, cannot improve on a direct mapping from phonemes to words. Lexical feedback might change the dynamics of processing, but it could not provide greater accuracy than a purely feedforward system, since it cannot, for example, improve the resolution of sublexical representations.

The second argument is that *interaction entails hallucination*. If interaction is permitted at low levels (i.e., perceptual levels), veridical perception by definition becomes impossible, since the information not present in the signal will be added. While one might argue that it would be possible to balance top-down and bottom-up sources of information, the typical rejoinder is that there is not a principled way to balance the two. From this perspective, attempts to quantify new sources and balances are ad-hoc and experiment-specific.

The last argument is against *radical* interaction – interaction among relatively distant representations, such as discourse and phonemes, or between modalities. The argument is that *context is infinite*, and unless interaction at low levels is prohibited or at least very tightly restricted, information processing will become intractable.

The latter two arguments are weak. Indeed, integrating top-down information will alter perception, but not randomly; as I will argue in a moment, it simply provides needed context dependence. Balancing top-down and bottom-up information is a matter of learning, for any given state of the system, the past relative reliability of sources of information. The same answer applies to the infinite context argument. We know multiple sources of information are integrated eventually (e.g., preceding context in the following sentence helps at least partially disambiguate “bank:” *John lost his money when the bank collapsed*); thus, this argument simply defers the inevitable. We must determine empirically what sorts of information are integrated and how early (potentially immediately in the previous example, but in the following case no biasing information aside from sense frequency is available except post-perceptually: *when the bank collapsed, John lost his money*).

The first argument is rather more difficult. It is tempting simply to appeal to evidence of top-down effects, such as the word superiority effect (i.e., that phonemes can be detected more quickly in words than nonwords; Rubin, Turvey, & Van Gelder, 1976), or phoneme restoration (context-dependent restoration of a phoneme replaced with noise or an ambiguous sound as a function of lexical or sentential context, e.g., Warren, 1970; Samuel, 1981, 1997). Norris et al. argue most such effects are post-perceptual, and feedforward explanations are possible for others. In addition, in simulations with the TRACE model (McClelland & Elman, 1986), the premier example of an interactive activation network where lexical feedback plays an instrumental role in the system's dynamics, Frauenfelder and Peeters (1998) reported that when they turned feedback off, half the words they tested were recognized more quickly than they were with feedback on. This would seem to suggest that feedback in TRACE is simply providing a mechanism for accounting for top-down effects, but plays no functional role in the model.

However, there are two gaps in this argument. The first is noise, which is no small concern when we are talking about virtually any aspect of perception. Noise is of particular importance in the case of spoken language, given the tremendous sources of variability (phonetic context, speaking rate, talker characteristics, acoustical environment, background noise, etc.). One of the original motivations for feedback in interactive activation models was to make them robust in noise (McClelland, Rumelhart & Hinton, 1986). The second gap is context dependence. Feedback provides a basis for context-sensitive processing. In the case of lexical feedback, words provide implicit, context-sensitive prior probabilities for strings of phonemes. Consider a case where a phoneme is mispronounced, obscured by noise, or is otherwise ambiguous, with the result that the bottom-up input is slightly biased toward a nonword ("I told the poss" rather than "I told the boss"). All parties agree that eventually, top-down knowledge will resolve the ambiguity in favor of the lexical bias, but Norris et al. argue that the lexical bias should not come into play at early perceptual levels.

However, it is telling that Norris et al. are willing to accommodate any top-down knowledge that can be incorporated into feedforward connections, such as diphone transitional probabilities – their proposed solution for accommodating findings that transitional probabilities from phone A to phone B influence the processing of both (e.g., Pitt & McQueen, 1998). This appears to violate the assertion that one cannot do better than well-tuned bottom-up acoustic-phoneme and phoneme-lexical mappings, as incorporating transitional probabilities permits a small degree of context dependence. Technically, this context dependence should not be required, but it does not violate the bottom-up principle so long as the feedforward units remain sublexical and unmodulated by lexical or other top-down knowledge *during perception*. But if diphones are allowed, why not triphones, or any n -phone that is helpful? As it turns out, the n -phones required to explain the entire range of findings similar to those of Pitt and McQueen are dynamic – the length of n varies from context-to-context, such that the appropriate context is approximately equal to word length (Magnuson, McMurray, Tanenhaus, & Aslin, 2003). So far, no one has proposed a mechanism for instantiating dynamic contexts without appealing to feedback.

Allowing context reduces the argument from a strong position ("top-down feedback does not benefit speech recognition; it can hinder it," Norris et al., 2000, p. 299) to an assertion that context can help, but you do not need feedback to provide it; anything that can be done with feedback can be done with a purely feedforward system (and that feedforward systems are less complex than feedback systems and so should be preferred). It is well-known that for any nondeterministic system, there exists a deterministic solution (Ullman & Hopcroft, 1969; also, Minsky & Pappert, 1969, provide a related proof that for any recurrent network, there exists a feedforward network that can provide identical input-output mappings for any finite period of time).

There are three hitches, however, if one wishes to assert that a feedforward system can do anything a feedback system can. First, feedforward solutions that are equivalent to feedback solutions (e.g., feedforward networks that are behaviorally-equivalent to recurrent networks; Rumelhart, Hinton, & Williams, 1986) have to be (sometimes several times) larger, because the network must be reduplicated for every desired time step of history (violating Norris et al.'s paraphrase of Occam: "never ... multiply entities unnecessarily"). This is minor, though, compared to the second and third problems, which I will discuss in detail in the next two sections. The second problem is that while the autonomous view predicts perceptual feedback should not exist, and that it should hinder recognition if it does, there is considerable evidence for interaction. The final problem is that it is virtually impossible to learn the appropriate feedforward mapping *without supervision* (being told the correct mapping) for any system that is not "easily invertible" (Friston, 2003) – that is, for any nondeterministic many-to-many mapping from signal to percept – which we shall see is true of language.

Evidence Against Interaction? Tanenhaus, Leiman, and Seidenberg (1979) and Swinney (1979), in very similar studies, famously demonstrated apparent evidence for staged processing of lexical access and context integration. Tanenhaus et al. presented spoken homophones (e.g., *rose*) in sentences biased towards different homophone senses (e.g., *they all rose* vs. *they all bought a rose*). The task for subjects was to name a visually presented word as quickly as possible. When visual probes were presented at the offset of the spoken homophone, priming was found for all senses (e.g., *flower* and *stand*). If the probe was delayed 200 msecs, only context-appropriate priming was found. This is consistent with staged processing, in which a word recognition system activates all form matches, and a later stage of processing selects the context-appropriate form (and this is how it is typically presented in textbooks: "Thus, context does have an effect on word meaning, but it exerts its influence only after all meanings have been briefly accessed," Goldstein, 2005, p. 356). However, the story is much more complicated.

First, even when the probe is presented at homophone offset, there is a trend towards greater priming of the context-appropriate sense. This is consistent with continuous rather than staged integration, if one assumes it takes some time for a detectable degree of integration to occur (cf. McClelland, 1987). To make this more

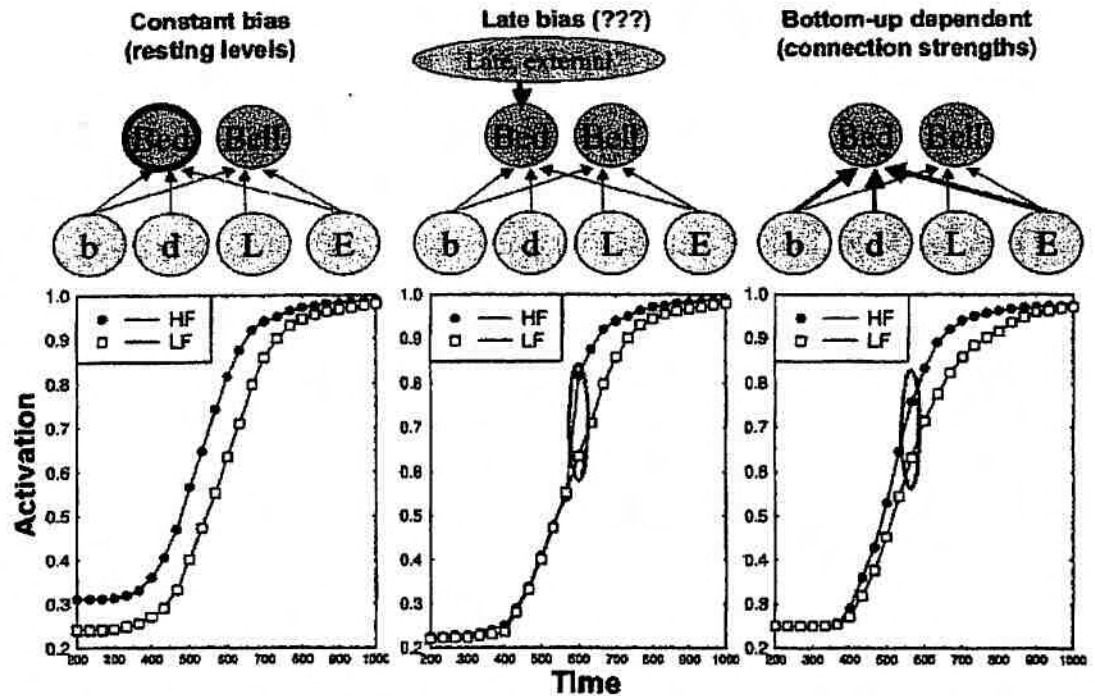


FIGURE 18.2 Schematics of three possible loci of frequency effects in spoken word recognition and predicted activation patterns.

concrete, I will use a parallel example having to do with the locus of word frequency effects. Conine, Titone, and Wang (1993) reported that frequency effects were not evident in fast responses (in a critical condition) to a lexical frequency biased stop continuum (e.g., *grass-crass*). This, among other results, led them to propose that frequency is a late-acting bias that is not an integral part of lexical representations. How else might we explain this effect?

Consider Figure 18.2, which contains schematics of three possible ways of implementing frequency in neural network models (limited for illustration to phoneme and lexical nodes). On the left, we have a constant bias, in the form of resting level bias (indicated by the bold circle around *bed*). Below the network, there is a plot comparing the resulting time course of activation for separate simulations with either *bed* (high frequency, or HF) or *bell* (LF) as the target. On this account, even if the system were forced to make a response early in the time course, there would always be a frequency effect. In the middle panels, the late-bias approach is schematized. Here, there is a constant frequency effect that kicks in at an unspecified point during processing (e.g., when lexical activations hit some threshold). If the system must respond prior to this "magical moment," no frequency effect will be observed. The oval indicates the point in time where the frequency effect can first be detected. A third possibility is shown in the right panels: making the bottom-up connection strengths proportional to frequency. The thicker lines emanating from phoneme nodes to *bed* (e.g., the thicker connection from /b/ to *bed* compared to *bell*) indicate stronger connections. Consistent with a general Hebbian

learning account, these connections are proposed to be stronger because they have been used more frequently. This arrangement leads to a substantially different activation pattern compared to the other two. There is a constant frequency effect, but it is proportional to the bottom-up input, so it starts out weak, and becomes stronger. The oval in this panel indicates a hypothetical moment when the difference becomes large enough to be detected in a button-press task like that used by Connine et al. Depending on the resolution of the measurement, a weak effect can masquerade as a late effect.

Dahan, Magnuson, and Tanenhaus (2001) revisited this finding using the “visual world” eye tracking paradigm. Participants saw a display with four objects in it. On critical trials, these were a low-frequency target (e.g., *bell*), a low-frequency competitor (e.g., *bench*), a high-frequency competitor (e.g., *bed*) and an unrelated distractor (e.g., *carriage*). We used the interactive activation model, TRACE (McClelland & Elman, 1986), to predict the proportion of fixations to each object at a fine-grained time scale. Consistent with a connection-strength implementation of frequency, we found early, continuous effects of frequency that increased over time (analogous to the pattern in the bottom right panel of Figure 18.1, but with our more complex design).

But what about the Tanenhaus et al. (1979) and Swinney (1979) results, which suggest there is an initially encapsulated stage of form access prior to the availability of larger context? There are two further wrinkles to consider. Several studies have demonstrated that how early one detects evidence of interaction in ambiguity resolution depends on (a) the strength of any bias that may exist between possible lexical interpretations (e.g., frequency differences in *rose=flower* vs. *rose=stood*; Simpson & Burgess, 1985) and (b) the strength of the contextual bias (Duffy, Morris, & Rayner, 1988). Consistent with these general findings, Shillcock and Bard (1993) revisited the homophone issue, focusing on a single pair: *would/wood*, for which there is a very strong frequency bias for *would*, which they augmented with strong contexts (*John said he couldn't do the job but his brother would* vs. *John said he couldn't do the job with his brother's wood*). In the modal-biased condition, they found no evidence of priming for *timber* when they probed at word offset, or even if they probed prior to word offset. Similarly, Magnuson, Tanenhaus, and Aslin (2002) found no evidence for activation of adjective-noun onset competitors (analogous to *purple-purse*) given strong pragmatic and syntactic expectations for one part of speech. Thus, it is eminently plausible to attribute the 1979 results to continuous interaction that is proportional to bottom-up input, and therefore is difficult to detect early in processing.

Evidence for interaction is not limited to adjacent levels of linguistic description. Another aspect of written and spoken language where strictly staged, autonomous models have dominated is sentence processing. The Garden Path Model (Frazier & Fodor, 1978; Frazier & Rayner, 1982) provides an elegant, compact theory of human sentence parsing, in which initial syntactic structure building proceeds without consideration of the semantics of individual words nor of any larger context. Instead, the parser proposes the simplest possible structure (following a small number of clearly operationalized heuristics, such as minimal attachment and late closure), which nearly always turns out to be correct. When it is not, reanalysis is required. Proponents of the Garden Path theory claimed it was possible

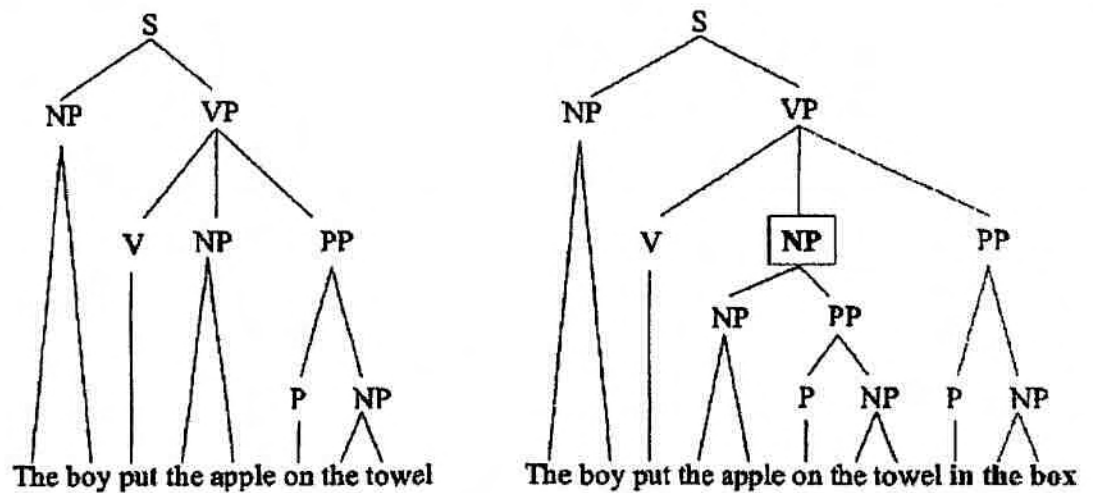


FIGURE 18.3 The simplest structure consistent with hearing *the boy put the apple on the towel* (left), and the more complex structure required to accommodate a sentence that continues *...in the box* (the boxed NP is the additional node required to parse *the apple on the towel* as a noun phrase rather than a noun phrase and location).

to account for apparent demonstrations of strong effects of lexical semantics and context (such as those reviewed by MacDonald, Pearlmutter & Seidenberg, 1994, and Trueswell & Tanenhaus, 1994), typically by proposing that reanalysis was able to operate too quickly to be detected reliably given strong contexts.

Figure 18.3 illustrates potential syntactic structures when the sentence being heard or read is "the boy put the apple on the towel in the box." By the time you get to "towel," the simplest structure is consistent with the towel being the goal location of "put." That structure is shown on the left in Figure 18.3. When you encounter "in the box," that structure no longer works. One of multiple possibilities is shown on the right in Figure 18.3. The structures explicitly handling "in the box" are in grey. Our focus is the new NP node (shown in a box) required to make "on the towel" modify "apple" (i.e., which apple – the one on the towel) rather than specify the goal location of "put." This new structure is more complex than the one on the left, and so is initially dispreferred. On this sort of model, though, it is crucial to note that the structure on the left should always be built first given a sentence like this one, without early reference to any sort of context.

Tanenhaus, Spivey-Knowlton, Eberhard and Sedivy (1995) tested a dramatically different interactive hypothesis. They provided potentially "unhelpful" and "helpful" visual contexts for instructions like, "put the apple on the towel in the box." Schematics of example contexts are shown in Figure 18.4. In the helpful case, there are two apples, so a complex referring expression is required to disambiguate between the two; "on the towel" specifies *which* apple. In the unhelpful case, there is only one apple, which means that the instruction, with respect to the display, violates Grice's (1975) maxim of quantity (do not be overly specific): there is no need to specify the location of the apple, since "the apple" would be unambiguous. In both contexts, there is an empty towel, allowing a possible goal interpretation of




FIGURE 18.4 Schematics illustrating the helpful and unhelpful visual displays from Tanenhaus et al. (1995).

"on the towel" in either case. On the Garden Path view, there should be an initial period of linguistic processing encapsulated from the visual display, and so the visual context should not have an early impact. On a highly interactive view, linguistic processing ought to be constrained by expectations governed by potential ambiguities in the display.

Tanenhaus et al. tracked eye movements as subjects followed the spoken instructions, and found differences between the two contexts from the earliest eye movements. In the helpful context condition, upon hearing "apple" subjects were equally likely to make a saccade to either apple, and on hearing "on" they quickly settled on the correct apple. They made no looks to the empty towel, and performed the expected action quickly. In contrast, given the unhelpful context, subjects made many looks to the empty towel, and were significantly slower to perform the expected action (and in some cases, subjects actually moved the apple from one towel to the other, or even picked up the apple and towel and moved both to the box). This is an example of unequivocally non-linguistic context influencing linguistic processing as early as we can measure, and so supports radical interaction.

Later studies using this technique have shown discourse, syntax and even lexical processing are moderated by visual contexts (e.g., Chambers, Tanenhaus, Eberhard, Filip, & Carlson, 2001; Sedivy, Tanenhaus, Chambers, & Carlson, 1999), affordances of held instruments (Chambers, Tanenhaus, & Magnuson, 2004), and even momentary changes in the affordances available to an interlocutor (Hanna & Tanenhaus, 2003).

But what of the finding that feedback in TRACE does not improve processing (Frauenfelder & Peeters, 1998)? To put it simply, this result has not held up. The original

simulations were conducted with a rather odd set of 21 words that were all seven phonemes long with a uniqueness point at the fourth phoneme. These items were selected for principled reasons for earlier simulations. We revisited this finding for two reasons: first, we questioned whether that set of 21 words would be representative of the lexicon, and second, the original simulations did not address the important motivation for feedback of making the system robust in noise (Magnuson, Strauss, & Harris, 2005). We tested a large lexicon (about 1000 words) with feedback on or off at multiple levels of noise. Without noise, 77% of the words were recognized more quickly with feedback on than with feedback off (many of the others showed no advantage, though a number were recognized much more quickly without feedback, which we attribute to peculiarities of their neighborhoods). When noise was added, the recognition time advantage persisted, and feedback preserved accurate recognition – the model was significantly less accurate without feedback when noise was added. (See also Mirman, McClelland, & Holt, 2005, for another case where TRACE's predictions have turned out to be correct despite previous claims that feedback in TRACE was inconsistent with empirical results).

Thus, contra arguments for autonomous levels of representation in language processing, there is evidence for interaction at any level we measure as early as we can measure. Furthermore, feedback provides benefits in simulations with TRACE, speeding processing (on average) and protecting the model against noise. But Norris et al. (2000) complain in their reply to the concurrently published commentaries on their article that no one has offered a theoretical case for the need for feedback. I present just such an argument in the next section, by appealing to principles of theoretical neuroscience.

Interaction, Ambiguity, and Representational Learning. In appealing to neurobiology for insight into the interaction question, we can begin by asking about the prevalence of feedback (reciprocal) connections in cortex. Not only are they plentiful, they outnumber forward connections, and span more cortical levels (Zeki & Shipp, 1988). So certainly there exist potential mechanisms for on-line interaction in the brain. Next, we can ask whether there are any clear cases where perceptual processes in any domain are mediated in an on-line fashion. The answer here is “yes,” from vision research, where the neural substrates are better understood than for language. For example, Rivadulla, Martinez, Varela, and Cudeiro (2002) report evidence that feedback connections modulate gain of thalamic receptive fields, Lee and Nguyen (2001) report that illusory contours activate V1 and V2 cells with timing consistent with feedback modulation, and Rao and Ballard (1999) argue that backward connections are required for a satisfactory account of dynamic context dependence of visual receptive fields. So there are precedents for interaction in perceptual-cognitive systems. How can we establish whether perceptual interaction underlies language processing, given that the neural basis for language remains rather sketchy? Here is where we must appeal to general learnability principles.

Again, while feedforward mappings (i.e., deterministic mappings) are possible for any nondeterministic system, the mapping cannot be learned without supervision

unless it is *easily invertible*. A system that is not easily invertible is one in which forward (cause to data) and inverse (data to cause) mappings are largely distinct. Another way of putting this is that nondeterministic systems (in which there are many-to-many mappings from causes to data, or stimuli to percepts) are not easily invertible because the effects of "causes" mix nonlinearly such that one cannot unambiguously recover the source of any given "datum." To approximate an unmixing in a feedforward network, the model must be given access to context-specific prior probabilities. That is, it requires a training signal – typically the expected or correct result – which means it requires explicit access to the outcome of the mapping, rather than opportunities to discover the mapping.

For illustration, Friston (2003) uses the example of $v = u^2$. Given v , it is impossible to know whether u is positive or negative. Such ambiguities are rampant in perception generally, and in language in particular, due to context dependence. Perceptual constancy depends upon mapping raw sensory stimulation to context-appropriate causes (e.g., attributing distinct wavelengths of light to the same color as a function of ambient illumination). In the case of speech perception, we experience phonetic constancy despite dependence on phonetic context (the same acoustic pattern signals a different consonant depending on the following vowel; Liberman, Cooper, Shankweiler, & Studdert-Kennedy, 1967), talker characteristics (two talkers' productions of different vowels may have identical formants, while their productions of the same vowel might have quite distinct formants [Peterson & Barney, 1952], and consonants are similarly talker-dependent [Dorman, Studdert-Kennedy, & Raphael, 1977]), and rate (the same acoustic pattern may be perceived as /b/ at one rate, but /w/ at another; Miller & Baer, 1983). And consider again the ambiguities in Figure 18.1, or global orthographic (*lead*), lexical (*bowl*), or sentential ambiguities (*the cop watched the spy with binoculars*). The point is that context dependency and ambiguity are *typical* of language at any level of analysis. This puts language in the domain of non-invertible mappings, and therefore, feedback is required to learn the context-dependent and ambiguous mappings required of language processing.

Friston (2003) reviews a variety of learning frameworks used in theoretical neuroscience, and compares them within an expectation maximization framework. He presents empirical Bayes as a neurobiologically plausible learning framework. This method does not require supervised learning. It does require a hierarchy of representations, however, and details of how the hierarchy itself is learned are sparse so far. We can appeal to mechanisms proposed in adaptive resonance theory (ART; e.g., Grossberg, 1986) for potential mechanisms to establish hierarchies. When a learning system is exposed to an input signal, it begins to compile recurring patterns into representational units (chunks). As patterns among chunks are discovered, compositional chunks are instantiated downstream of the smaller chunks. Empirical Bayes is a bootstrap method in which representations at level x provide estimated prior probabilities for level $x-1$, providing the basis for both forward and inverse models of the input, which is what is required to approximate a nondeterministic (non-invertible) mapping. Even though the representations at level x may be weak or noisy, their ability to provide additional constraint on the potential causes of the bottom-up signal is leveraged into distinct forward and inverse inferences (see

Friston, 2003, for technical details). Other frameworks have potential as well (indeed, adaptive resonance itself may be approximately equivalent), but the key point here is that feedback serves two purposes: it provides necessary information (prior probability estimates) for representational learning, and provides context-dependence for on-line processing.

Norris, McQueen, and Cutler (2003) recently reported evidence for rapid retuning of speech perception. To account for this learning without appealing to on-line feedback, they point out the logical difference between on-line feedback and what they call "feedback for learning" (e.g., the distinct step in typical connectionist models of backpropagating error). They agree feedback for learning is necessary, but deny this has any implications for their arguments against on-line feedback. They open the door to admitting evidence for on-line feedback, but only with "spandrel" status (a "spandrel" being an architectural feature that is not functionally necessary but results from combinations of necessary features, such as the spaces between joined gothic arches): on-line feedback may exist, but only because mechanisms are required for feedback for learning, and, for unknown reasons, those feedback mechanisms may operate continuously – providing useless on-line feedback as well as necessary feedback for learning. However, while spandrels in evolutionary biology typically refer to exaptations (putting the spandrel to use, such as a bird using its wings to shield its eyes, or a snail making use of its umbilicus [a groove formed incidentally in shell formation] as a brooding chamber; Gould & Lewontin, 1979), Norris et al. maintain that even if evidence of on-line feedback is found, its on-line use *logically cannot provide any benefit to language processing*.

The autonomous theory is significantly weakened by this position (cf. McClelland, Mirman, & Holt, 2006). It hedges its bets by allowing that on-line feedback may exist, but rather than offering a theoretical explanation for the large and varied number of results that support interaction reviewed in the previous section, it predicts that on-line feedback should not exist, and if it does, the theory claims it cannot serve any useful purpose and in fact should hinder speech recognition. Interactive theories explicitly predict effects of on-line feedback as well as feedback for learning, hold that feedback provides benefits (speed, accuracy, and protection against noise), and are consistent with biological evidence for the prevalence of feedback connections and their beneficial use in vision (Lee & Nguyen, 2001; Rao & Ballard, 1999; Rivadulla et al., 2002). As I have reviewed here, language falls into a class of representational learning problems that require feedback for learning (since the forward and inverse models are distinct; Friston, 2003), and for which feedback is eminently useful in on-line processing (in helping resolve nondeterminancies via context-specific prior probabilities of dynamic size).

The implication, therefore, is that a focus on any single level (such as words) must be considered provisional in two respects. First, the arguments in this section present a strong case that language is highly interactive, which limits the inferences that should be drawn when observation of the system is limited to a discrete level such as the single word. Second, as I discussed under the first concern (division of labor), the focus must be recognized as a functional theoretical assumption, capable of masking useful information available to the listener/reader because of the

hypotheses it leads us to consider, as well as those it dissuades us from considering. When language is viewed developmentally, both of these concerns should be amplified, as the relative autonomy of a single level may change over development, as a function of development within the language system (cf. Seymour, this volume), as well as in other cognitive and social domains.

PART 2: THE PSYCHOLINGUIST AS SISYPHUS

Will we be able to run any psycholinguistic experiments at all in 1990? (Anne Cutler, 1981)

Every year it seems that new constraints on language processing are discovered. At the lexical level, there is a growing list of factors to be controlled that includes frequency, neighborhood, uniqueness point, word length, phonotactic probability, prosodic context – and on and on. Just when one might assume a debate is settled – e.g., that semantics influences single word naming (Strain, Patterson, & Seidenberg, 1995) – someone appeals to the growing list of lexical characteristics to find one that may provide an alternative explanation (in this case, age of acquisition; Monaghan & Ellis, 2002). Debate ensues as to the degree to which various measures can be considered independent, etc. (Strain, Patterson, & Seidenberg, 2002), and rolling a boulder up a hill for all eternity starts sounding not so bad by comparison. The ever-expanding list of factors that must be controlled led Anne Cutler to the (perhaps only somewhat) tongue-in-cheek title quoted at the beginning of this section. Two tools that can provide considerable leverage on this “material dilemma” are artificial materials and computational models.

Let's Make Stuff Up

In addition to the sheer number of lexical characteristics that have been identified, we must grapple with the fact that establishing independent influences of these factors is made very difficult because (a) subsets of them tend to be highly correlated, and (b) the degree to which they are based on strong theoretical principles varies (age of acquisition and imageability can be operationalized in a lexicon-external fashion, but metrics like neighborhood cannot, and their theoretical interpretation depends crucially on assumptions regarding lexical activation and competition). Words in natural languages do not fall into neat strata that would allow easy factorial exploration of lexical characteristics, and the facts that the characteristics tend to correlate and vary in theoretical transparency makes even regression approaches less than compelling.

Kessler, Treiman, and Mullennix (this volume) suggests a way out of this dilemma would be to train subjects on artificial linguistic materials, so as to have precise control over the materials. An additional advantage would be the fact that, in principle, individual differences in factors such as vocabulary would be somewhat mitigated. Artificial orthographies were used productively over the last few decades mainly to study acquisition processes (e.g., Byrne & Carroll, 1989; Knafle & Legenza, 1978), and this approach has been extended recently to examine neural

effects of alphabeticity and training conditions (Bitan & Karni, 2003, 2004; Bitan, Manor, Morocz, & Karni, 2005). More to the point for the 'material dilemma', this method has more recently been applied to issues of control. Two impressive uses of this approach are described in this volume in chapters by Mauer and McCandliss and Hart and Perfetti. While their results are compelling, their methods are rather daunting, as they require considerable training. For example, Mauer and McCandliss review a study by McCandliss, Posner, and Givon (1997) in which subjects were trained for 50 hours. Such a study obviously requires tremendous amounts of labor, analysis, and motivated subjects. While one can address exceedingly complex questions with this much training (and one can go quite a bit further; pairs of subjects in a study by Hudson & Eigsti [2003] spent 30 hours over 9 weeks learning 233 Farsi words and using those to communicate with each other, which allowed Hudson & Eigsti to observe analogs to Pidginization processes in the lab, as pairs developed various syntactic mechanisms), one can do quite a lot with much less training.

For example, my colleagues and I used artificial spoken materials to set up factorial comparisons of different sorts of phonological overlap, frequency, and neighborhood density (Magnuson, Tanenhaus, Aslin, & Dahan, 2003). Subjects learned 16 words that referred to novel geometric forms. After two days of training (with about 1 hour per day), performance resembled that with real words, and we were also able to track the emergence of competition dynamics as new neighborhoods of artificial words were learned.

With similar amounts of training, we were also able to address the cross-form class competition questions I discussed earlier (Magnuson et al., 2002). We wanted to test whether items like *purse* and *purple* compete when pragmatics lead to strong expectations for a noun or adjective. For example, consider Figure 18.5. In the top panel, where there are two purses and two cups (one of which is meant to be purple), if I ask you, "hand me the pur-," *purple* should be a strong pragmatic candidate at that moment, quite possibly stronger than *purse*. Given the display on the bottom, *purse* is a much better candidate than *purple*, as a simple noun is much more likely, since it would be sufficient for unambiguous reference. A crucial question regarding interaction is *when* those pragmatic influences kick in – do they constrain initial lexical activation, or do they apply post-perceptually? However, the kinds of pairs we would need, like *purple-purse*, are relatively hard to come by, and are hopelessly uncontrolled for phonological overlap, frequency, neighborhood, etc. (e.g., *tan-tambourine*, *dotted-dog*, *rough-rug*). By using an artificial lexicon where "nouns" referred to geometric shapes and "adjectives" referred to textures applied to the shapes, we had precise control over not just lexical characteristics, but also pragmatic expectations, as we determined the sets of items in the display and the regularity with which conversational norms were obeyed. We found that pragmatic constraints had immediate impact; when the display predicted adjective use, there was no competition between adjectives and nouns, and vice-versa.

Two days of training is still a lot, and impractical for many populations (such as children). We are currently exploring simplified and more engaging versions of this paradigm for use with children. However, there is a new study closer to our topic of single word reading that resolves a rather slippery debate with a design limited to one-hour, including training and testing. Trudeau (2006) re-examined the so-called



FIGURE 18.5 Conditions that require an adjective (or other complex referring expression) for unambiguous reference (top), versus conditions where an adjective is not required (bottom).

Strain effect (Strain et al., 1995). In the original study, an interaction of regularity and imageability was found. This would support the “triangle” model (Harm & Seidenberg 2004; Plaut, McClelland, Seidenberg, & Patterson, 1996; Seidenberg & McClelland, 1989) over the dual route model (Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001), since the former predicts an integral, low-level role for semantics in word recognition while the latter does not. Monaghan and Ellis (2002) argued that imageability was confounded with age of acquisition. Trudeau recognized this as a perfect opportunity for using artificial materials. He assembled a set of 60 nonwords based on real English words with irregular pronunciation (e.g., BINT and MAVE). Subjects learned high- or low-imageability definitions for the “new” words, as well as their pronunciation. Two lists were used with different groups of subjects. Each item was trained with a regular pronunciation in one list (BINT rhymes with MINT) and irregular in the other (BINT rhymes with PINT – though note that pronunciation training was based on actual audio recordings of the pronunciations, not analogies to real words). Within the space of an hour, subjects learned the items to an 80% criterion and were given a rapid naming test. Trudeau observed an interaction of regularity and imageability similar to that found by Strain et al. (1995).

Caution is warranted with artificial language materials, however. The largest cause for concern is whether results with artificial materials will generalize to natural language. Frequency effects, for example, are implemented in artificial language studies by manipulating number of exposures. But it is not apparent that relatively short-term repetition frequencies are comparable to the long-term prior probabilities that underlie frequency effects with natural language materials. Another concern is that the artificial materials may interact with native language knowledge.

While there is evidence that artificial materials are largely functionally isolated from the native lexicon in typical studies (Magnuson et al., 2003, Experiment 3), the likelihood of interference will depend on tasks and the nature of the artificial materials (indeed, Ramscar [2002] exploited interaction with the native lexicon to examine past-tense inflection of novel words presented as though they were new words in English). To state these concerns broadly, effects observed under the idealized conditions of artificial language studies may not generalize back to natural language processing.

A solution is to treat artificial language studies as exercises in prototyping. With knowledge in hand of how the factors studied in an artificial language experiment interact, one can refine hypotheses of what to expect with natural language materials – and if possible, those hypotheses should be tested, rather than trusting only in the artificial language results. On the other hand, artificial language studies can be extremely helpful when results with natural materials are ambiguous or in dispute (as in Trudeau's [2006] replication of the Strain effect, i.e., the regularity by imageability interaction).

Where Are the Models?

Theories of language processing continue to grow more complex, reflecting the amazingly rich database of empirical results. Once a theory includes even a handful of interacting theoretical assumptions, the behavior of an implemented system with properties corresponding to the assumptions becomes difficult to predict, even analytically, let alone based on box-and-arrow diagrams. When a theory arrives at a fairly low level of complexity, *simulation* in an implemented model becomes the best (if not perhaps the only reliable) method of predicting the actual behavior that results from a set of theoretical assumptions. There have been several examples in spoken word recognition where quite logical inferences about what various theories or models would predict were suggested in the literature and generally accepted, only to turn out to be completely wrong when someone finally attempted to verify the predictions with an implemented model (see Magnuson, Mirman, & Harris, in press, for a review).

So it is surprising to see so little modeling work in the domain of single word reading. On the one hand, there are the familiar debates between proponents of the dual route model (Coltheart et al., 2001) and the triangle model (Harm & Seidenberg, 2004; Plaut et al., 1996; Seidenberg & McClelland, 1989), which arguably provide similar coverage for basic effects in visual word recognition, as well as for different types of dyslexia (though the mechanisms underlying the explanations are often quite dissimilar; for recent overviews, see Coltheart, 2005, and Plaut, 2005). However, the triangle model has important advantages compared to the dual route model. First, it is a learning model, which makes it amenable to testing predictions of how reading should change over time as the component skills of reading develop (Harm, McCandliss, & Seidenberg, 2003; Powell, Plaut, & Funnell, 2006), and even to testing predictions regarding the efficacy of different interventions (Harm et al., 2003). Recent work by Harm and Seidenberg (2001, 2004) has set a new standard of

specificity, analyzing how cooperative and competitive elements of the model lead to efficient reading. Their 2004 paper includes an ambitious attempt to incorporate semantics into a model. Harm and Seidenberg found that the model predicts that early in training, phonological representations provided the main pathway to accessing meaning, but with expertise, the balance shifted to include more direct access from orthography (though both pathways remained instrumental).

This last finding is generally consistent with discoveries that part of the left fusiform gyrus is engaged strongly in the processing of print (Cohen et al., 2000). Three somewhat different takes on the organization of this "visual word form area" (and the areas and circuits involved in reading more generally) appear in the chapters in this volume (Frost et al.; Mauer & McCandliss; Nazir & Husckauf; Simos et al.). The behavioral and neuroimaging evidence appear to depend on expertise, tasks (including precise control and presentation of visual stimuli; Nazir & Husckauf), materials, and the transparency of subjects' native language orthography (Maurer & McCandliss). The comprehensive approaches represented in this volume present tremendous challenges for theory development. Theories must synthesize behavioral and neuroimaging studies of normal and disabled readers of various ages, reading native language or artificial language materials (and note that the VWFA may be analogous to some degree to the iconographic procedure within Seymour's theory [Seymour, this volume]). Model simulations provide partial means for bootstrapping theories from these disparate sources of data by identifying principles that govern not just visual expertise in reading, but its weight relative to phonological skill as reading develops in different populations. Models also provide a tool somewhat like artificial language material studies. One can set up idealized cases and examine what predictions emerge from the model. Then, one can observe which factors lead to significant changes in model performance as more realistic conditions and more factors are added. In particular, this approach could shed light on complex effects like the feedback consistency results described by Kessler et al. (this volume). Another case where modeling could provide true insights is in comparing the three approaches to spelling development reviewed by Pollo, Treisman, and Kessler (this volume). Pollo et al. suggest that the distinct accounts favored primarily in English-speaking countries (the phonological perspective that the child must grasp the alphabetic principle before she can master spelling) and Romance-speaking countries (the constructivist perspective, which predicts specific stages that appear not to take place in English spelling development) may be the result of distinct pressures on spelling development in the different languages. On their view, a general statistical learning account might explain both patterns as a function of language-specific features. This is an ideal question for modeling: can one model provide both developmental patterns as a function of language-specific features?

On the other hand, the range of data reported in this volume may also provide the means for distinguishing between competing models. For example, it is not clear whether current models could capture hypothesized changes in phonological and orthographic grain size as reading develops (Goswami, this volume; though one might be able to test whether an analog to grain size can be detected in hidden unit space as a connectionist model like the triangle model develops). Also, the advent of great specificity in describing brain regions and circuits, in typical and

poor readers, young and old (e.g., Frost et al.) suggests this area is ripe for symbiotic model use, where models may be able to guide interpretation of neural data, and the neural data may provide the basis for pushing models from Marr's (1982) *algorithmic* level to a level somewhat closer to his *implementational* level. But for models to guide theories, or for results to test models, the models must be used.

In this volume, only one chapter makes significant reference to model predictions. Grainger and Ziegler, in their chapter on cross-code consistency, explain a series of results with reference to figures diagramming an interactive activation framework. Their explanations are compelling, and I do not doubt that actual simulations with an implemented model could come out as they predict. Indeed, their work on intermodality neighborhood effects is truly groundbreaking (Ziegler, Muneaux, & Grainger, 2003), and their model schematics provide a convincing account of the effects. However, this falls into the domain of what Kello and Plaut (2003) call "fundamentalist" rather than "realist" use of modeling. That is, the hypothetical model provides a framework for describing a set of theoretical principles. In this case, the predictions, while complex, follow so sufficiently transparently from the model diagrams that actual simulations may not be needed. Even if they were, though, they would provide a *fundamentalist* demonstration of the general principles of cross-code consistency. What would not be clear is whether they would hold up in a *realist* model: one with a large vocabulary, and the ability to provide broad coverage of the empirical findings on reading beyond cross-code consistency (and again, there is always the risk that actual simulations might diverge from the predictions).

So why is so little modeling done? One reason is that modeling is technically challenging. Even when existing models are made publicly available (e.g., the dual route model is available at: <http://www.maccs.mq.edu.au/~max/DRC/>), they tend to require significant computer and/or programming skill to use or modify. If model developers wish to see modeling adopted widely, one strategy would be to develop user-friendly versions, ideally with graphical user interfaces, that would make modeling approachable for the average researcher (cf. "jTRACE," a cross-platform, easy-to-use and modify implementation of the TRACE model [McClelland & Elman, 1986] of speech perception and spoken word recognition; Strauss, Harris, & Magnuson, in press).

CONCLUSIONS

I have addressed two distinct themes. The second was the simpler of the two: I reviewed two tools (artificial language material studies and computational models) that can be used to "prototype" experiments by setting up very clear tests of hypotheses. In the case of artificial language materials, the researcher has control over experience with the materials to be tested, and all their properties (orthographic transparency and neighborhood, phonological neighborhood, frequency, etc.). This allows one to design studies in which dimensions hypothesized to be important (e.g., imageability and regularity; Trudeau, 2006) can be manipulated more strongly than is often possible with natural materials, and without potential interference from the numerous confounding and extraneous variables at play in typical psycholinguistic

studies. Computational models can serve a similar function via idealized simulations where only the dimensions of interest vary, but also provide the means to examine how multiple theoretical assumptions will interact. Lately, models of reading have provided not just accounts of typical effects in adult reading (frequency, neighborhood, length, etc.) but also of reading development and even remediation (Harm et al., 2003; Harm & Seidenberg, 2004). I also bemoaned the absence of modeling in most work on single word reading, and advocated the development of implemented models that are easy enough to use for the average researcher.

The more complex theme was theoretical. I argued that while studies focused on single word reading have been extremely productive and provide a necessary window on many component skills of reading, as well as vital diagnostics for reading disability, caution is warranted when focusing on single words. Identifying any level of description as a discrete level of representation may overestimate the complexity of the information processing problem (by masking potentially useful bottom-up and top-down constraints) and overestimate modularity of language processing; even if one assumes discrete levels as a simplifying assumption, the questions asked and the explanations considered tend to be influenced by those assumptions. I also provided a theoretical and empirical case for the argument that language is highly interactive. The interactive nature of the language processing system implies that the focus on single words must be provisional, and the goal in developing theories of reading must be to integrate word-level theories into larger theories of comprehension.

ACKNOWLEDGMENTS

Discussions with Harlan Harris, Ted Strauss, and Dan Mirman fostered many of the ideas presented here, and Inge-Marie Eigsti's insightful comments improved this chapter considerably. Preparation of this chapter was supported by NIH grants DC005765 to James Magnuson, and HD001994 and HD40353 to Haskins Laboratories.

REFERENCES

- Abramson, M., & Goldinger, S. (1997). What the reader's eye tells the mind's ear: Silent reading activates inner speech. *Perception & Psychophysics*, *59*, 1059–1068.
- Andruski, J. E., Blumstein, S. E. & Burton, M. (1994). The effect of subphonetic differences on lexical access. *Cognition*, *52*, 163–187.
- Ashby, W. R. (1962). Principles of the self-organizing system. In von Foerster, H. & G. W. Zopf, Jr. (Eds.), *Principles of Self-organization*, pp. 255–278. Pergamon Press.
- Bitan, T. & Karni, A. (2003). Alphabetical knowledge from whole words training: effects of explicit instruction and implicit experience on learning script segmentation. *Cognitive Brain Research*, *16*, 323–337.
- Bitan, T. & Karni, A. (2004). Procedural and declarative knowledge of word recognition and letter decoding in reading an artificial script. *Cognitive Brain Research*, *19*, 229–243.

- Bitan, T., Manor, D., Morocz, I. A., & Karni, A. (2005). Effects of alphabeticality, practice and type of instruction on reading an artificial script: An fMRI study. *Cognitive Brain Research*, *25*, 90–106.
- Byrne, B., & Carroll, M. (1989). Learning artificial orthographies: Further evidence of a non-analytic acquisition procedure. *Memory & Cognition*, *17*(3), 311–317.
- Chambers, C. G., Magnuson, J. S., & Tanenhaus, M. K. (2004). Actions and affordances in syntactic ambiguity resolution. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *30*, 687–696.
- Chambers, C.G., Tanenhaus, M.K., Eberhard, K.M., Filip, H., & Carlson, G.N. (2001). Circumscribing referential domains in real-time language comprehension. *Journal of Memory and Language*, *47*, 30–49.
- Cohen, L., Dehaene, S., Naccache, L., Lehericy, S., Dehaene-Lambertz, G., Henaff, M. (2000) The visual word form area: Spatial and temporal characterization of an initial stage of reading in normal subjects and posterior split-brain patients. *Brain*, *123*, 291–307.
- Coltheart, M. (2005). Modeling reading: The dual-route approach. In M. J. Snowling & C. Hulme (Eds.), *The Science of Reading: A Handbook*, pp. 6–23. Oxford: Blackwell.
- Coltheart, M., Rastle, K., Perry, C., Langdon, R. & Ziegler, J. (2001). DRC: A Dual Route Cascaded model of visual word recognition and reading aloud. *Psychological Review*, *108*, 204–256.
- Connine, C. M., Titone, D., & Wang, J. (1993). Auditory word recognition: Extrinsic and intrinsic effects of word frequency. *Journal of Experimental Psychology: Learning Memory and Cognition*, *19*(1), 81–94.
- Cutler, A. (1981). Making up materials is a confounded nuisance: or, Will we be able to run any psycholinguistic experiments at all in 1990? *Cognition*, *10*, 65–70.
- Dahan, D., Magnuson, J. S., and Tanenhaus, M. K. (2001). Time course of frequency effects in spoken-word recognition: Evidence from eye movements. *Cognitive Psychology*, *42*, 317–367.
- Dahan, D., Magnuson, J. S., Tanenhaus, M. K., and Hogan, E. M. (2001). Tracking the time course of subcategorical mismatches: Evidence for lexical competition. *Language and Cognitive Processes*, *16* (5/6), 507–534.
- Davis, M. H., Marslen-Wilson, W. D. & Gaskell. M. G. (2002). Leading up the lexical garden-path: segmentation and ambiguity in spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 218–244.
- Dorman, M. F., Studdert-Kennedy, M., & Raphael, L. J. (1977). Stop consonant recognition: Release bursts and formant transitions as functionally equivalent, context-dependent cues. *Perception & Psychophysics*, *22*, 109–122.
- Duffy, S. A., Morris, R. K., Rayner, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory and Language*, *27*, 429–446.
- Fodor, J. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Frazier, L. & Fodor, J.D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, *6*, 1–34.
- Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, *14*, 178–210.
- Frauenfelder, U. H. & Peeters, G. (1998). Simulating the time course of spoken word recognition: an analysis of lexical competition in TRACE. In J. Grainger and A. M. Jacobs (Eds.), *Localist connectionist approaches to human cognition* (pp. 101–146). Mahwah, NJ: Erlbaum.
- Friston, K. (2003). Learning and inference in the brain. *Neural Networks*, *16*, 1325–1352.
- Goldstein, E. B. (2005). *Cognitive Psychology: Connecting Mind, Research, and Everyday Experience*. Belmont, CA: Thomson Wadsworth.

- Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London, B* 205, 581–598.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and Semantics, Vol. 3, Speech Acts*, pp. 41–58. New York: Academic Press.
- Grossberg, S. (1986). The adaptive self-organization of serial order in behavior: Speech, language, and motor control. In: E.C. Schwab & H.C. Nusbaum (Eds.), *Pattern recognition by humans and machines, Vol. 1, Speech perception*, pp. 187–294, Orlando: Academic Press, Inc.
- Hanna, J.E. & Tanenhaus, M.K. (2003). Pragmatic effects on reference resolution in a collaborative task: evidence from eye movements. *Cognitive Science*, 28, 105–115.
- Harm, M., McCandliss, B.D., & Seidenberg, M.S. (2003). Modeling the successes and failures of interventions for disabled readers. *Scientific Studies of Reading*, 7, 155–182.
- Harm, M., & Seidenberg, M.S. (2001). Are there orthographic impairments in phonological dyslexia? *Cognitive Neuropsychology*, 18, 71–92.
- Harm, M., & Seidenberg, M.S. (2004). Computing the meanings of words in reading: cooperative division of labor between visual and phonological processes. *Psychological Review*, 111, 662–720.
- Hudson, C. L., & Eigsti, I. M. (2003). The Lexical Competence Hypothesis: A cognitive account of the relationship between vernacularization and grammatical expansion in creolization. *Journal of Pidgin and Creole Languages*, 18(1), 1–79.
- Kahre, J. (2002). *The Mathematical Theory of Information*. Boston: Kluwer Academic Publishers.
- Kello, C. T., & Plaut, D. C. (2003). Strategic control over rate of processing in word reading: A computational investigation of the tempo-naming task. *Journal of Memory and Language*, 48, 207–232.
- Knafle, J. D., & Legenza, A. (1978). External generalizability of inquiry involving artificial orthography. *American Educational Research Journal*, 15(2), 331–347.
- Landi, N., Perfetti, C. A., Bolger, D. J., Dunlap, S., & Foorman, B. R. (2006) The role of discourse context in developing word representations: A paradoxical relation between reading and learning. *Journal of Experimental Child Psychology*, 94(2), 114–133.
- Lee, T. S. & Nguyen, M. (2001). Dynamics of subjective contour formation in early visual cortex. *Proceedings of the National Academy of Sciences*, 98, 1907–1977.
- Lehiste, I. (1972). The timing of utterances and linguistic boundaries. *Journal of the Acoustical Society of America*, 51, 2018–2024.
- Lieberman, A.M., Cooper, F. S., Shankweiler, D. P., and Studdert-Kennedy, M. (1967). Perception of the speech code. *Psychological Review*, 74, 431–461.
- Lukatela, G., Eaton, T., Sabadini, L., & Turvey, M. T. (2004). Vowel duration affects visual word identification: Evidence that the mediating phonology is phonetically informed. *Journal of Experimental Psychology: Human Perception and Performance*, 30, 151–162.
- Lukatela, G., & Turvey, M. T. (1994a). Visual lexical access is initially phonological: 1. Evidence from associative priming by words, homophones, and pseudohomophones. *Journal of Experimental Psychology: General*, 123, 107–128.
- Lukatela, G., & Turvey, M. T. (1994b). Visual lexical access is initially phonological: 2. Evidence from phonological priming by homophones and pseudohomophones. *Journal of Experimental Psychology: General*, 123, 331–353.
- MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review*, 101(4), 676–703.
- Magnuson, J. S., McMurray, B., Tanenhaus, M. K., and Aslin, R. N. (2003). Lexical effects on compensation for coarticulation: A tale of two systems? *Cognitive Science*, 27, 795–799.
- Magnuson, J. S., Mirman, D., & Harris, H. D. (in press). Computational models of spoken word recognition. In M. Spivey, M. Joanisse, & K. McRae (Eds.), *The Cambridge Handbook of Psycholinguistics*. Cambridge University Press.

- Magnuson, J. S., Strauss, T. J. & Harris, H. D. (2005). Interaction in spoken word recognition models: Feedback helps. *Proceedings of the XXVII Annual Conference of the Cognitive Science Society*, 1379–1384. Lawrence Erlbaum Associates.
- Magnuson, J. S., Tanenhaus, M. K. and Aslin, R. N. (2002). Immediate integration of syntactic and referential constraints on spoken word recognition. In W. D. Gray & C. Schunn (Eds.), *Proceedings of the 24th Annual Conference of the Cognitive Science Society*, pp. 614–619. Mahway, NJ: Erlbaum.
- Magnuson, J. S., Tanenhaus, M. K., Aslin, R. N., and Dahan, D. (2003). The microstructure of spoken word recognition: Studies with artificial lexicons. *Journal of Experimental Psychology: General*, 132(2), 202–227.
- Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman.
- Marslen-Wilson, W., & Warren, P. (1994). Levels of Perceptual Representation and Process in Lexical Access: Words Phonemes and Features. *Psychological Review*, 101, 653–675.
- McCandliss, B. D., Posner, M. I., & Givon, T. (1997). Brain plasticity in learning visual words. *Cognitive Psychology*, 33, 88–110.
- McClelland, J. L. (1987). The case for interactions in language processing. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading*, pp. 3–36. Erlbaum.
- McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18, 1–86.
- McClelland, J. L., Mirman, D., & Holt, L. L. (2006). Are there interactive processes in speech perception? *Trends in Cognitive Sciences*, 10(8), 363–369.
- McClelland, J. L., Rumelhart, D. E. & Hinton, G. E. (1986). The appeal of parallel distributed processing. In D. E. Rumelhart, J. L. McClelland, and the PDP research group. *Parallel distributed processing: Explorations in the microstructure of cognition. Volume I* (pp. 3–44). Cambridge, MA: MIT Press.
- McQueen, J.M., Cutler, A., Briscoe, T. & Norris, D. (1995). Models of continuous speech recognition and the contents of the vocabulary. *Language and Cognitive Processes*, 10, 309–331.
- Miller, J. L., & Baer, T. (1983). Some effects of speaking rate on the production of /b/ and /w/. *Journal of the Acoustical Society of America*, 73, 1751–1755.
- Minsky, M. L. & Papert, S. A. (1969). *Perceptrons*. Cambridge: MIT Press.
- Mirman, D., McClelland, J. L., & Holt, L. L. (2005). Computational and behavioral investigations of lexically induced delays in phoneme recognition. *Journal of Memory and Language*, 52, 424–443.
- Monaghan, J., & Ellis, A.W. (2002). What, exactly, interacts with spelling-sound consistency in word naming? *Journal of Experimental Psychology: Learning, Memory & Cognition*, 28, 183–206.
- Nicholson, T. (1991). Do children read words better in context or in lists? A classic study revisited. *Journal of Educational Psychology*, 83, 444–450.
- Norris, D., McQueen, J.M. & Cutler, A. (2000). Merging information in speech recognition: Feedback is never necessary. *Behavioral and Brain Sciences*, 23, 299–325.
- Norris, D., McQueen, J.M., & Cutler, A. (2003). Perceptual learning in speech. *Cognitive Psychology*, 47, 204–238.
- Peterson, G. E. & Barney, H. L. (1952). Control methods used in a study of vowels. *Journal of the Acoustical Society of America*, 24, 175–184.
- Peterson, G., & Lehiste, I. (1960). Durations of syllabic nuclei in English. *Journal of the Acoustical Society of America*, 32, 693–703.
- Pitt, M. A., & McQueen, J. M. (1998). Is compensation for coarticulation mediated by the lexicon? *Journal of Memory and Language*, 39, 347–370.
- Plaut, D. C. (2005). Connectionist approaches to reading. In M. J. Snowling & C. Hulme (Eds.), *The Science of Reading: A Handbook*, pp. 24–38. Oxford: Blackwell.

- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*, 56–115.
- Port, R. (1981). Linguistic timing factors in combination. *Journal of the Acoustical Society of America*, *69*, 262–274.
- Powell, D., Plaut, D. C., and Funnell, E. (2006). Does the PMSP connectionist model of single word reading learn to read in the same way as a child? *Journal of Research in Reading*, *29*, 299–250.
- Rao, R. P., & Ballard, D. H. (1998). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive field effects. *Nature Neuroscience*, *2*, 79–87.
- Rivadulla, C., Martinez, L. M., Varela, C., & Cudeiro, J. (2002). Completing the corticofugal loop: A visual role for the corticogeniculate type 1 metabotropic glutamate receptor. *Journal of Neuroscience*, *22*, 2956–2962.
- Rubin, P., Turvey, M. T., & van Gelder, P. (1976). Initial phonemes are detected faster in spoken words than in spoken nonwords. *Perception & Psychophysics*, *19*, 384–398.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel and Distributed Processing: Exploration in the Microstructure of Cognition, Vol. 1*, pp. 318–362. MIT Press, Cambridge, Massachusetts.
- Salverda, A. P., Dahan, D., & McQueen, J. M. (2003). The role of prosodic boundaries in the resolution of lexical embedding in speech comprehension. *Cognition*, *90*, 51–89.
- Samuel, A. G. (1981). Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General*, *110*(4), 474–494.
- Samuel, A. G. (1997). Lexical activation produces potent phonemic percepts. *Cognitive Psychology*, *32*(2), 97–127.
- Sedivy, J.C., Tanenhaus, M.K., Chambers, C., & Carlson, G.N. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition*, *71*, 109–147.
- Seidenberg, M. S. and McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*, 523–568.
- Shillcock, R. C. & E. G. Bard. (1993). Modularity and the processing of closed class words. In Altmann, G.T.M. & Shillcock, R.C. (Eds.) *Cognitive models of speech processing. The Second Sperlonga Meeting*, pp. 163–185. Erlbaum.
- Simpson, G. B., & Burgess, C. (1985). Activation and solution processes in the recognition of ambiguous words. *Journal of Experimental Psychology: Human Perception and Performance*, *11*, 28–39.
- Strain, E., Patterson, K., & Seidenberg, M. S. (1995). Semantic effects in single-word naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition* *21*, 1140–1154.
- Strain, E., Patterson, K., Seidenberg, M. S. (2002). Theories of word naming interact with spelling-sound consistency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 207–241.
- Strauss, T. J., Harris, H. D., & Magnuson, J. S. (in press). jTRACE : A reimplementaion and extension of the TRACE model of speech perception and spoken word recognition. *Behavior Research Methods, Instruments & Computers*.
- Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, *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 Behavior*, *18*, 427–440.
- Tanenhaus, M.K., Spivey-Knowlton, M.J., Eberhard, K.M. & Sedivy, J.E. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*, 1632–1634.
- Trudeau, J. (2006). *Semantic Contributions to Word Naming with Artificial Lexicons*. Unpublished Ph.D. dissertation, University of Connecticut Department of Psychology.

- Trueswell, J.C. & Tanenhaus, M. K. (1994). Toward a lexicalist framework of constraint-based syntactic ambiguity resolution. In Clifton, and Frazier (Eds), *Perspectives on sentence processing*, pp. 155–179. Hillsdale, NJ: Lawrence Erlbaum.
- Ullman, J. & J. E. Hopcroft. (1969). *Formal Languages and Their Relation to Automata*. Reading, MA: Addison Wesley.
- Warren, R. M. (1970). Restoration of missing speech sounds. *Science*, 167, 392–393.
- Zeki, S., & Shipp, S. (1988). The functional logic of cortical connections. *Nature*, 335, 311–317.
- Ziegler, J. C, Muneau, M. & Grainger, J. (2003). Neighborhood effects in auditory word recognition: Phonological competition and orthographic facilitation. *Journal of Memory and Language*, 48, 779–793.