Connectionist models of language processing and the training of listening skills with the aid of multimedia software

Hulstijn, J.H.

Link to publication

Connectionist Models of Language Processing and the Training of Listening Skills With the Aid of Multimedia Software

Jan H. Hulstijn
University of Amsterdam, Amsterdam, The Netherlands

ABSTRACT

This paper gives a brief overview of recent developments in cognitive science, juxtaposing two schools of thought, symbolism and connectionism. It is argued that both connectionism and symbolism are relevant for second language learning. A simple software tool for listening comprehension is then described, illustrating how some principles of connectionism as well as symbolism can be implemented in computer-assisted language learning.

1. INTRODUCTION

Most literate, educated people believe that when they listen to the news and hear the utterance ''The doctors found a small increase in blood pressure'', what they actually hear is a sentence, consisting of phrases, consisting of words, consisting of sounds corresponding to letters. They believe, in other words, that they hear discrete units. This is not the case however. Let us take the word 'small' in the context of 'found a small increase' as an example. Why didn't they hear 'foundism all-in crease', 'found is mall increase', 'dismal increase', 'us mall', or other, even less meaningful partitions of the incoming sounds? What reached their ears was a wave of sounds without borderlines. The word 'sound' may already be misleading because, physically, each 'speech sound' is, as it were, (1) a blend of itself and the previous sound, and (2) a blend of itself and the following sound. Thus, the ears do not 'hear' discrete speech sounds. What the brain of skilled language users appears to be doing is transforming input

consisting of meaningless tiny bits of acoustic information into meaningful units, such as speech sounds, syllables and words. How the brain exactly performs this amazing feat is still not well understood in detail (Carroll, 2002; Gernsbacher, 1994) but in what follows it is argued that becoming a good listener is, in essence, a matter of building a huge network of brain cells capable of computing probabilities of possible 'language-like solutions', such as 'is mall', 'dismall', 'foundism all-in', and 'a small' and many other, less or non-meaningful solutions. Section 2 of this paper elaborates the view that language acquisition in general, and developing proficiency and fluency in parsing and understanding speech in particular, is a matter of constructing a neural network. Section 3 presents the implications of such a view for instructional programs that aim to help learners of a second language (L2) to gain fluency in L2 listening with the aid of multimedia technology.

2. SYMBOLISM AND CONNECTIONISM

Cognitive science is the name of a multidisciplinary enterprise in which philosophers, biologists, psychologists, neuroscientists, linguists and other scholars participate to disentangle the mysteries of the human mind in relation to the mysteries of the brain. The binding term is cognition. Cognition is used as an umbrella term to cover human capabilities to perceive information, to encode and store perceived information, to retrieve information, to think (e.g., in order to solve problems), and to behave skillfully on the basis of stored information. One of the goals of cognitive science is to develop computational models (so called architectures) of the representation and processing of knowledge.

In many respects, human behavior can be said to be regular. For instance, when we are playing a game of chess, we act according to rules. Similarly, when we produce a linguistic utterance, the utterance has a structure, which can be described with grammar rules. Does the fact that many aspects of our behavior – in chess, in verbal communication, and in many other walks of life, – to the extent that it is regular, can be described with rules, mean that we actually base behavior on the application of mental rules? It is to this question that symbolists and connectionists give different answers (Pinker, 1999).[1]

---

[1]The literature on this issue is sheer endless. Pinker (1999) offers a nontechnical, reader friendly, but not unbiased, rendition of the main issues.

The symbolist school of thought claims that our cognition consists of abstract categories, and rules that operate on these categories. Thus, in the domain of language, linguistic cognition can best be described with a grammar, containing categories such as S, VP, V, NP, DET and N, and by rules, such as S → NP VP, VP → V NP, NP → DET N (S). The grammar also has a lexicon, containing lexical items, that is, form-meaning connections provided with category labels, for example, *the* (DET), *man* (N), and *kiss* (V). Grammars need few types of tools (basically only categories and rules) to be able to not only account for a large number of productions (sentences), but also for the infinity of language (with the rules above, sentences can be produced such as *The man kissed his wife, who kissed their daughter, who kissed her brother, who . . .*), and, even more importantly, for the regularity of these productions. In sum, symbolism deals with regularity and irregularity in two fundamentally different ways: regular phenomena are accounted for by means of rules and categories, irregular phenomena are listed in the lexicon.

The connectionist school of thought takes a radically different approach to cognition. The aim of connectionists is to build architectures capable of doing three things: account for (1) the representation of knowledge, (2) the acquisition (and loss) of knowledge, and (3) the use of knowledge. A connectionist architecture or system is a network of interrelated units or nodes, representing knowledge. The connection between any two nodes is said to have a certain activation weight, reflecting the strength with which the two nodes are associated. When the system is exposed to new information (in the form of input nodes), the connection between some of its nodes may increase or decrease somewhat (acquisition or loss of knowledge). When asked to perform (use of knowledge), the system produces output nodes which reflect its current internode activation patterns.

There are several types of connectionist architectures (Dijkstra & De Smedt, 1996). The two main classes are (1) localist networks, which consist of interconnected symbolic, meaningful categories (such as phonemes, syllables, codas and word stems), and (2) parallel distributed processing (PDP) networks, which consist of subsymbolic, nonmeaningful nodes. For instance, in a PDP model of the recognition of written words, letters are recognized on the basis of letter features such as straight/curved, horizontal/vertical/diagonal lines; and words are recognized on the basis of trigrams (Rumelhart & McClelland, 1986). The word *small*, for instance, preceded and followed by a space (#), is represented by the combination of the following patterns: *#sm, sma, mal, all*, and *al#*. Thus, in a PDP network, there is no single node

representing the word *small*; instead, the representation of the word *small* is distributed over five (meaningless) trigrams.

Acquisition is modeled via an activation flow that runs through the network. When a word-recognition network which takes letters as its input, is exposed to the stimulus word *small*, the activation of the five trigram nodes mentioned above is raised a little, while the activation of other trigrams (e.g., *smo*, *sno*, *tma*, *tna*) is simultaneously lowered (inhibition). When many texts are fed to the system, word by word, it will gradually develop relatively high activation values for patterns of trigrams that form possible words (e.g., *small*, *mall*, *smell*, *mill*, *slim*) and relatively low activation values for patterns that form infrequent or nonexistent words (e.g., *smal*, *smla*, *sllim*). After a training session involving a subset of the words of English, including *small*, *mall*, *smell*, *mill* and *slim*, but excluding the word *slam*, we can expose the system, in a subsequent testing session, to the letter strings *slam* and *lsam*, and observe how much activation each word will receive. There is sufficient evidence that PDP networks are capable of responding with a high activation pattern in the case of *slam* but with a low activation pattern in the case of *lsam*. This is another way of saying that the system has accepted the new string *slam* as (relatively) grammatical and rejected *lsam* as relatively ungrammatical. We could thus say that the system has 'learned' some knowledge on the basis of which it can deal with information to which it has never been exposed before. When, during training, the system has also been fed with stimuli like *smaller*, *taller*, *bigger*, and *older*, and many other comparative forms of adjectives, it will accept *smarter* even when this word was not a member of the training set.

The crucial difference between a subsymbolic PDP network and a symbolic grammar is that the former treats words like *bigger*, *older* and *smarter* just as associated patterns of meaningless units, such as trigrams, whereas the latter treats them as consisting of two categories, a word stem and a comparative suffix *er*, combined by a rule. Thus, both types of architectures are capable of dealing with a number of phenomena of regularity, but they do so in different ways. Furthermore, a PDP network deals with regular and irregular forms in identical ways (by representing both types as patterns of subsymbolic units), whereas a symbolic grammar deals with regular and irregular phenomena separately (by creating rules for regular forms and storing irregular forms in the lexicon).

What can we conclude and learn from the debate between symbolists and connectionists? In the course of this debate, which has been going on for almost 20 years now, some issues of contest have been resolved and many

others are not likely to be resolved soon. However, as new theoretical claims and new improvements of the architectures have been put forward, the debate helps us to get a clearer understanding of the secrets of human cognition in general and of knowledge and use of language in particular. In the following three subsections, some of these issues, relevant in the context of computer-assisted L2 instruction, will be briefly mentioned.

2.1. The Mind-Brain Issue

It is not easy to assess the value of the arguments for and against various types of architectures because it is often left implicit at which point of the mind-brain continuum a certain architecture is proposed to have explanatory value. Is the architecture supposed to have a philosophical/mental, a psychological/behavioral, or a neurobiological/neurophysiological function? A generative grammar can best be seen as an attempt to explain mental phenomena, at the mind end of the continuum. Because of its serial nature (the generation of a sentence is a stepwise, serial, non-parallel procedure), a grammar is better suited to account for the representation of linguistic knowledge, what Chomsky (1986) called competence, than for the on-line processing of it. For instance, a generative grammar cannot adequately reflect the speed with which we process linguistic information on-line, during listening and speaking. Connectionist architectures of the so called localist type, consisting of networks of symbols (such as phonemes, codas, syllables, word stems, word endings, etc.) and allowing processes to take place in parallel, have been primarily developed to account for behavioral data, in particular for the speed and accuracy of human language use. For instance, localist models of speaking aim to account not only for accurate speech but also for occasional speech errors (Levelt, Roelofs, & Meyer, 1999). If the model produces an error, the error should be a 'human error'. For example, it should allow for the production of the occasional human error *a pig bark* instead of the intended *a big park*, but not for the implausible error *a bag pirk*.

Many proponents of PDP models have claimed both psychological and neurophysiological plausibility for these architectures. The neurophysiological plausibility of PDP networks has been said to reside in the resemblance of the subsymbolic nature of knowledge representation in PDP networks with the way neurons and their axons and synapses in brain tissue are interconnected, and the resemblance of the mechanics of activation spreading in PDP networks with the way in which electrochemical processes take place in the brain, involving the secretion and diffusion of neurotransmitters. However,

others have argued that there are substantial differences between PDP networks and the brain and that the resemblance is only superficial without much significance (McLeod, Plunkett, & Rolls, 1998, 8–15).

2.2. The Necessity of Hybrid Architectures

An increasing number of scholars now believe that both connectionist and rule-based architectures are needed to account for two types of cognition which humans appear to posses: on the one hand the ability to associate great numbers of stimuli with each other and distill fuzzy, prototype-like categories from them (which can best be modeled with associative network archi-tectures), and, on the other hand, the ability to produce complex, regular, mental productions in a large variety of human activities, such as problem solving, engineering, developing theories, and, forming linguistic utterances. It is likely that for many regular behavioral phenomena, humans possess two coexisting forms of cognition, one of a network type and another of the rule type. For instance, when confronted with a new but regular stimulus, such as *traller*, both systems begin to process the information; either or both systems may or may not be successful in producing an adequate output. But if both do succeed, one may be quicker than the other. In that case, the behavioral response is said to be based on the system which won the race. This notion of what is often called a 'horse race' between two competing systems is now quite common in psycholinguistics (e.g., Baayen & Schreuder, 2002; MacWhinney, 1989; Rayner & Pollatsek, 1989).

2.3. Item-Based Learning and the Driving Force of Frequency

It may well be that the acquisition of many forms of cognition, including language, takes place in several phases: (1) the accumulation of a number of information units (often referred to as items, instances, or exemplars), (2) the building of a network of these units with different association strengths between them, reflecting frequency and regularity effects of the input, and (3) eventually the forming of abstract categories and combinatorial rules. This acquisition order may be at work simultaneously at different levels and domains. For instance, at some stage, language learners may be simulta-neously working on the accumulation of phonological, morphonological, lexical and syntactic knowledge, such that, in each domain, (1) some knowl-edge has the form of an associated network, (2) in which some common patterns are emerging, (3) some of which will result in the construction of rules.

The acquisition of common patterns and the construction of categories is driven by frequency (Ellis, 2002). It is through exposure to large amounts of input that learners implicitly learn that some patterns are highly common while others are not. For instance, the phoneme combination *sm* is common in English whereas **ms* is not; and it is through exposure to large amounts of language that native speakers of English implicitly learn that verbs like *give* and *offer* differ from verbs like *donate* and *present* because the former category allows dative alternation whereas the latter does not: whereas *he gave a present to his sister* and *he gave his sister a present* are both correct, only *she donated some money to the university* is correct but **she donated the university some money* is not.

3. SOME IMPLICATIONS FOR COMPUTER-ASSISTED LANGUAGE LEARNING

We have seen, in the previous section, that the acquisition of cognition is, in terms of formal architectures, first of all a matter of the accumulation of information units in a network and adjusting the activation values of inter-unit connections on the basis of frequency of stimuli in the input. This is true for L2 acquisition too. This means, in practice, that L2 learners must get the opportunity to make themselves familiar with the phonetic and phonological properties of the L2, learn large amounts of words, and automatize their ability to recognize words in speech. Human beings have a limited capacity to pay conscious attention to information. Our attention capacity is too small to pay conscious attention to several things simultaneously. Yet we are capable of doing several things at the same time, provided that most of these things take place automatically, not demanding conscious control. When we read, processes of word recognition normally take place automatically. In contrast, children, in initial stages of learning to read, pay so much attention to reading individual words that they have no attention capacity left for the meaning of what they read; when reading the fourth word of a sentence they may have already forgotten the first word. Fluent reading and listening is characterized by automatic processing at the lower levels of word recognition and sentence parsing, leaving attention capacity free to concentrate on the higher levels of information, that is, on semantics and content (Harrington, 2001; Rost, 2002).

I have argued elsewhere (Hulstijn, 2001, 2002, in press) that vocabulary acquisition and the automatization of word-by-word understanding processes

may be the two most crucial factors to boost L2 acquisition. Furthermore, there are good reasons to regard listening as the most implicit and least explicit one of the four language skills, more implicit indeed than reading (see also Harrington, 2001; Rost, 2002). Whereas readers often have the possibility to stop, think, and reread a sentence, listeners are forced to process speech in the same speed with which it is produced by the speaker. Whereas readers can clearly see the borderlines between words, as they are marked by interword spaces, listeners are exposed to concatenated words; they must acquire complex skills to segment speech into words, which requires a context-sensitive representation of phonemes and phoneme clusters both within and across word boundaries. Thus, listening to concatenated speech is even more difficult than reading a news trailer on an illuminated bulletin board without interword spaces, because phonemes, in speech, have less constancy than letters have in print, which was illustrated with the 'found a small increase' example in the introduction section of this paper. Furthermore, whereas readers can pause in order to consult and apply any explicit language knowledge they may possess, listeners have no time to do so (except in the case of speech intermitted with pauses). Speech segmentation hardly lends itself to conscious monitoring and needs to be acquired by extensive practice. Moreover, speech segmentation processes in L1 play an interfering role in L2 speech segmentation because of their automatic and implicit nature. Thus, L1 interference in word recognition may be harder to overcome in listening than in reading.

We may conclude from these facts that the acquisition of word recognition skills in listening requires special attention in the L2 curriculum. Can the computer be effectively used in the acquisition of these skills? Yes it certainly can, as I will try to demonstrate in the following section. Yet, a major breakthrough in the history of L2 pedagogy was *not* the advent of the computer and modern, digital multimedia technology, but rather the invention, more than 100 years ago, of the phonograph and similar devices with which sound could be recorded, stored, played, and replayed. Until then it was not possible for individual L2 learners to listen to a recorded stretch of speech as often as they needed in order to be able to fully understand what was spoken – except when they could afford a private native speaker/instructor. Thus, in principle, modern multimedia software is not needed to give learners the opportunity to listen again and again to a recording. A simple audio cassette player along with the printed text suffice, plus a piece of paper with which the printed text can be temporarily covered. With these simple tools, learners

(1) listen to the recording, (2) think for themselves whether they have understood everything what was said, (3) replay the recording as often as they deem necessary, (4) reveal the text by lowering the piece of paper to read what they have just heard, (5) realize what they should have understood, (6) replay the recording as many times as is necessary to be able to understand everything what was spoken without the aid of the written text.

4. A SOFTWARE TOOL FOR THE TRAINING
OF WORD-BY-WORD UNDERSTANDING

In Section 2 it was argued that it is essential to give L2 learners' cognitive systems the opportunity to accumulate and categorize acoustic, phonemic, syllabic, morphological and lexical information. In Section 3, a six-step procedure was described to provide learners with the necessary practice, using conventional tools. With the aid of multimedia software, this procedure lends itself to a much more convenient implementation. This section describes a simple piece of software, called 123LISTEN, designed at the University of Amsterdam to perform just this task.[2]

   The procedure to produce software for word-by-word understanding with 123LISTEN, consists of three steps. First, the language instructor selects a piece of video (or audio) of not more than several minutes (approximately 400 words). The selection must be based on curricular criteria determined by the goals of the language course: Which kinds of discourse are the learners supposed to be able to understand? (see, for instance, the list of discourse types mentioned in The Common European Framework of Reference for Languages, 2001, Chapter 4). Secondly, the materials should ideally promote learners' learning motivation. Furthermore, during the search for appropriate texts, instructors must bear in mind the requirement that only spoken texts which contain few unfamiliar words are suitable to train word recognition skills. The simple reason for this is that learners cannot 're-cognize' words which they do not yet 'cognize' to some extent. If a text contains words, likely to be unknown to the students, a L1 translation of these could be provided in parentheses in the speech transcript (see below).

---

[2]123LISTEN was developed by Jan de Jong (Faculty of Humanities, University of Amsterdam). More information about the application and how to obtain a demo can be found at ⟨www.hum. uva.nl/dsp/dido/listen⟩ or e-mail to J.R.deJong@uva.nl

The second step in the production of courseware consists of the digitalization of the selected video or audio text (if it does not already have a digital format) and the copying of the material into 123LISTEN.[3] The third step consists of text editing. In the editing mode of 123LISTEN, the language instructor can play the recording as many times as is necessary in order to divide the recording up into fragments of several seconds each. A fragment normally consists of one utterance, preferably marked at its beginning and end by a pause in the speech stream. The editing mode enables the instructor to mark the borderlines between fragments in terms of frame numbers (which are displayed on the screen). This is done with simple mouse clicks. The result can then be listened to; if the instructor is not satisfied, the borderline between fragments can be adjusted by clicking on forward ($>>$) or backward ($<<$) buttons. The next thing to be done is to type in the transcript of the speech to be heard in the fragment. When the beginning and end frames of all fragments have been marked and the transcripts have been entered (and edited, if transcription errors have been made), the text is ready for use by the students.

In the student mode, 123LISTEN first asks the student to select a text, and then provides three ways of listening:

(1) Nonstop playing of the text while listening and watching, without text. This is like normal television watching or like normal listening to the radio (if 123LISTEN is used with audio recordings).
(2) Playing the text fragment-by-fragment with delayed text display.
(3) Playing the text fragment-by-fragment with simultaneous text display. This is like watching a movie with subtitles in the language being spoken (like the facility often offered to hearing-disabled people).

The true purpose and real value of 123LISTEN is to be found in the second mode, that is, in playing the text fragment-by-fragment, with delayed display

---

[3]Many video and audio materials do exist already in digitalized form, whether purchased in stores (movies, documentaries) or downloaded from internet. Thus, instructors may like to consider whether such materials meet their requirements. Also, some of these readily available materials go with the printed, or printable, text of what is being spoken. This saves the language instructor the time of making transcriptions. In most countries it is not illegal to take small sections of TV or radio programs and use them, within the school or university, for study purposes by registered students. Often it suffices to send a letter to the broadcasting company telling that the school will use a 4-min section of a particular program for a language course, that the course materials will not be sold as a commercial product, and that the school is willing to pay a small fee to the possessor of the rights of the TV or radio program.

of the subtitles. Before they begin to work with 123LISTEN for the first time, learners should receive information concerning the rationale of mode 2 and be encouraged to use it before (or even without) using mode 3. Mode 2 allows for the sequence of steps, described at the end of Section 3, in the use of traditional audio cassettes and hard-copy transcripts.[4] This means that students working in mode 2 of 123LISTEN, (1) listen as often as they want to each fragment, by clicking on the Replay button, (2) try to reconstruct the utterance being spoken for themselves, (3) verify their prediction by clicking on the Text button and reading the subtitles, and (4) click on the Replay button as often as they wish to make sure they understand every word, for instance when they had misunderstood or missed some words.[5]

5. CONCLUDING REMARKS

Since their invention, computers have been used in L2 instruction mainly for the application of explicit, declarative and rule-based knowledge (vocabulary and grammar drills). In later years, CD ROMs were produced providing learners with comprehensive language courses, comprising the presentation and explanation of new data (e.g., a dialogue plus vocabulary and grammar explanations) and exercises in which the new elements had to be used. The internet now offers L2 learners opportunities to get large amounts of L2 input of their liking. However, more is needed than substantial amounts of exposure to make L2 instructional programs successful. The construction of an implicit knowledge base, allowing for fast and effortless access to linguistic knowledge in all four language skills, also requires the deliberate training of automatic word-processing skills. The computer offers an efficient medium with which such skills can be trained, especially in the areas of listening and reading. Various types of reading tasks can be devised which force learners to increase their reading speed (Coady, 1997; Hulstijn, 2001). More importantly, listening skills must be trained, as has been argued and illustrated in this paper.

---

[4]Empirical research in the relative effectiveness of mode 2, in comparison to modes 1 and 3, is underway at the University of Amsterdam.

[5]The use of 123LISTEN, as proposed here, presupposes that learners are literate (in L2). Software to help illiterate L2 learners develop lower-order listening skills, would require that each fragment is presented twice, first in its normal concatenated form, and later, upon the learner's request, in a rendition with short pauses between words. This would require that the text be recorded twice, as normal, concatenated speech and as speech with interword pauses.

It is a widespread belief among language instructors (and learners), that it suffices to help learners develop top-down inferencing skills in coping with oral and written texts, beyond their current linguistic competence. However, this is only half of what instructors need to do. The other, and perhaps more essential half of any successful language program should consist of tasks in which learners can automatize their bottom-up processing of linguistic information. The more learners are able to process text without effort at the lower levels of word recognition and sentence parsing, the more attention capacity is available for the processing of information at the higher levels of meaning and content. In other words, and returning to the metaphor of computational architectures introduced at the beginning of this paper, we may conclude that modern multimedia software tools offer excellent means to help learners build associative networks allowing for fast, parallel processing of linguistic information.

REFERENCES

Baayen, R.H., & Schreuder, R. (2002). The balance of storage and computation in the menta lexicon: The case of morphological processing in language comprehension. In S. Nooteboom, F. Weerman, & F. Wijnen (Eds.), *Storage and computation in the language faculty* (pp. 61–92). Dordrecht, The Netherlands: Kluwer.

Carroll, S.E. (2002). Induction in a modular learner. *Second Language Research*, *18*, 224–249.

Chomsky, N. (1986). *Knowledge of language: Its nature, origin and use*. New York: Praeger.

Coady, J. (1997). L2 vocabulary acquisition through extensive reading. In J. Coady & T. Huckin (Eds.), *Second language vocabulary acquisition* (pp. 225–237). Cambridge: Cambridge University Press.

Common European Framework of Reference for Languages: Learning, teaching, assessment. [Council for Cultural Co-operation, Education Committee, Modern Languages Division, Strasbourg]. Cambridge: Cambridge University Press, 2001.

Dijkstra, T., & De Smedt, K. (1996). Computer models in psycholinguistics: An introduction. In T. Dijkstra & K. De Smedt (Eds.), *Computational psycholinguistics* (pp. 3–23). London: Taylor and Francis.

Ellis, N. (2002). Frequency effects in language processing: A review with implications for theories of implicit and explicit language acquisition. *Studies in Second Language Acquisition*, *24*, 143–188.

Fodor, J.A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Gernsbacher, M.A. (Ed.). (1994). *Handbook of psycholinguistics*. San Diego: Academic Press.

Harrington, M. (2001). L2 sentence processing. In P. Robinson (Ed.), *Cognition and second language instruction* (pp. 91–124). Cambridge: Cambridge University Press.

Hulstijn, J.H. (2001). Intentional and incidental second-language vocabulary learning: A reappraisal of elaboration, rehearsal and automaticity. In P. Robinson (Ed.), *Cognition*

and second language instruction (pp. 258–286). Cambridge: Cambridge University Press.

Hulstijn, J.H. (2002). Towards a unified account of the representation, processing, and acquisition of second-language knowledge. *Second Language Research*, *18*, 191–221.

Hulstijn, J.H. (in press). Psycholinguistic perspectives on second language acquisition. In J. Cummins & C. Davison (Eds.), *Kluwer handbook on English language teaching*, Vol. II: *Classroom practice and teaching methods in ELT: Current trends and issues.* Dordrecht: Kluwer.

Levelt, W.J.M., Roelofs, A., & Meyer, A.S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, *22*, 1–75.

MacWhinney, B. (1989). Competition and connectionism. In B. MacWhinney & E. Bates (Eds.), *The crosslinguistic study of sentence processing* (pp. 422–457). Cambridge: Cambridge University Press.

McLeod, P., Plunkett, K., & Rolls, E.T. (1998). *Introduction to connectionist modeling of cognitive processes*. Oxford, UK: Oxford University Press.

Pinker, S. (1999). *Words and rules: The ingredients of language*. London: Weidenfeld & Nicolson.

Rayner, K., & Pollatsek, A. (1989). *The psychology of reading*. Englewood Cliffs, NJ: Prentice-Hall.

Rost, M. (2002). *Teaching and researching listening*. Harlow, UK: Longman.

Rumelhart, D.E., & McClelland, J.L. (1986). On learning the past tenses of English verbs. In J.L. McClelland, D.E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition*. Vol. 2: *Psychological and biological models* (pp. 272–326). Cambridge, MA: MIT Press.