The architecture of the bilingual word recognition system: From identification to decision*

TON DIJKSTRA
WALTER J.B. VAN HEUVEN
NICI, University of Nijmegen

The paper opens with an evaluation of the BIA model of bilingual word recognition in the light of recent empirical evidence. After pointing out problems and omissions, a new model, called the BIA+, is proposed. Structurally, this new model extends the old one by adding phonological and semantic lexical representations to the available orthographic ones, and assigns a different role to the so-called language nodes. Furthermore, it makes a distinction between the effects of non-linguistic context (such as instruction and stimulus list composition) and linguistic context (such as the semantic and syntactic effects of sentence context), based on a distinction between the word identification system itself and a task/ decision system that regulates control. At the end of the paper, the generalizability of the BIA+ model to different tasks and modalities is discussed.

1 Introduction

In the last decade, there has been a strong increase in the number of experimental studies investigating aspects of bilingualism, in particular the bilingual mental lexicon. There are various reasons for this upsurge of interest in the process of bilingual word recognition and the system that takes care of it. For instance, it raises issues that are not present in monolinguals, such as how bilinguals distinguish words from the two languages while reading or listening. Furthermore, if monolingual and bilingual word recognition reflects a common underlying processing system, bilingual research offers new ways of testing hypotheses derived from the monolingual domain. At the same time, an analysis of the bilingual word recognition system may enhance our understanding of its monolingual counterpart. In fact, if one assumes that recognition of words in different languages is subserved by one system, the question arises whether word recognition in one language can be appropriately studied apart from other-language knowledge. Yet another reason for investigating bilingual word recognition is that there are practical

consequences to this research for educational purposes (e.g., how to better teach a foreign language in the classroom). Finally, because words are the basic building blocks of sentences, it is important to understand how words are recognized in a bilingual context as a prerequisite for understanding how sentences are parsed by bilinguals. For all these research issues, reaching a detailed understanding of the bilingual word recognition system is an important aim.

In 1998, we presented a model for bilingual word recognition that was based on the empirical evidence available at the time (Dijkstra and Van Heuven, 1998; Dijkstra, Van Heuven and Grainger, 1998(a); Van Heuven, Dijkstra and Grainger, 1998; see also Grainger and Dijkstra, 1992). The model, called the Bilingual Interactive Activation (BIA) model, was concerned with the recognition of orthographic representations only. Nevertheless, because it was implemented on the computer, it allowed a precise simulation of the results of a series of experimental studies. Now, several years and many studies later, it is time to evaluate the model and to explore how it may be updated and extended in the light of recent evidence.

In the following, we will first summarize the BIA model and discuss the major empirical phenomena that it can account for. Next, we will discuss problems of the BIA model arising from limitations in its lexical and language representations, its handling of context effects, and its lacking an implemented task

Address for correspondence
Ton Dijkstra, NICI, University of Nijmegen, P.O. Box 9104, 6500 HE Nijmegen, The Netherlands.
E-mail: *dijkstra@nici.kun.nl*

structure. We will propose solutions to these problems that can be tested and implemented in future versions of the model. These solutions imply a major change in the BIA model with respect to the language nodes, as well as the addition of representations and a task/decision component. Because the new model incorporates, to a large extent, the old BIA model, we will call it the BIA+ model.

The BIA+ model makes a clear-cut distinction between an encapsulated word identification system and a task/decision system, which accounts for a larger and more diverse set of empirical findings. Linguistic information from input signal or (sentence) context may affect the word identification system, while non-linguistic context information (e.g., participants' expectations and strategies) influences parameter settings in the task/decision system. The new model assumes interactivity within the word identification system and between this system and higher-order systems such as the parser. However, the model also proposes that the lexical activation levels within the word identification system itself are not affected by the task/decision system and, therefore, not by sources of non-linguistic information either.

Apart from considering how words in stimulus lists and sentence contexts are recognized, we will consider to what extent assumptions of the basic model can be generalized to bilingual auditory word recognition and word production, and how a participant's performance is affected by the task at hand.

2 The Bilingual Interactive Activation (BIA) model

Visual word recognition in the monolingual domain has been defined as the retrieval of orthographic representations from the mental lexicon corresponding to the input letter string (Grainger and Dijkstra, 1996). It can be defined likewise in the bilingual domain. However, two theoretical issues arise in bilingual word recognition that do not apply in the monolingual case. First, during monolingual word recognition, word candidates become activated that are similar to the input string (Coltheart, Davelaar, Jonasson and Besner, 1977; Andrews, 1989; Grainger and Segui, 1990). With respect to bilingual word recognition, one can ask if lexical candidates from both languages are activated in bilingual word recognition, or from one language only. In the empirical literature, these two positions with respect to bilingual processing have been referred to as the language non-selective versus the language selective access hypotheses.

A related second issue is whether the lexical representations of their two languages are stored together in an integrated lexicon (irrespective of the language they belong to) or in different lexicons, separate for each language. Basically, in an integrated lexical system competition or selection effects may occur between lexical candidates of both languages, whereas in two separate lexical systems competition effects are limited to candidates of one language only.

This distinction between structural and processing aspects of bilingual word recognition is not often made in the bilingual literature. Generally, two rather than four theoretical view points have been contrasted: language selective access in independent lexicons versus language non-selective access to an integrated lexicon. The Bilingual Interactive Activation (BIA) model defends the second viewpoint by assuming that lexical access is basically non-selective in nature and that the bilingual mental lexicon is integrated across languages.

The BIA model, illustrated in Figure 1, shares the basic architecture and parameter settings of the monolingual Interactive Activation model (McClelland and Rumelhart, 1981). Apart from introducing an integrated Dutch and English lexicon, a major extension is the addition of a representational layer containing two "language nodes" that are connected to all the word nodes in both lexicons. The model implements top-down language-to-word inhibition (see Figure 1). Throughout this article, two crucial issues will be (a) whether the activity of the language nodes is influenced by non-linguistic information sources from outside the word recognition system, and (b) whether inside the system the language nodes significantly affect the activation (and therefore the recognition) of words presented in stimulus lists or sentences. In the BIA model, we argued "yes" to both questions. For the BIA+ model, we will revise our claims and argue "no" to both (but we will further argue that syntactic and semantic aspects of sentence context *can* affect the recognition of a following target word).

When a string of letters is presented to the BIA model, this visual input affects particular features at each letter position, which subsequently excite letters that contain these features and at the same time inhibit letters for which the features are absent. The activated letters next excite words in both languages in which the activated letter occurs at the position in question, while all other words are inhibited. At the word level, all words inhibit each other, irrespective of the language to which they belong. Activated word nodes from the same language send activation on to the corresponding language node, while activated language nodes send inhibitory feedback to all word

which language a particular word belongs. The need for this type of representation has already been motivated decades ago (e.g., McCormack, 1977, p. 63). Second, language nodes collect activation from the lexical representations within a language. A carry-over of such *global lexical activation* from one trial to the next could account for between-trial language priming effects (e.g., Grainger and O'Regan, 1992) if one assumes that language nodes activated by an item can affect word activation at the next trial (top-down effect), or if decision criteria are affected by the same/different language membership of consecutive items (bottom-up effect).[1] Third, the language nodes serve as a functional mechanism modulating relative language activation to account for performance differences across experiments. Thus, language nodes can function as a sort of *language filter* (rather than an all-or-none language switch). Fourth, it might be assumed that the language nodes can collect context activation originating from outside the word recognition system. Through *contextual pre-activation*, the language nodes can account for potential top-down effects on the lexical identification system originating from, e.g., the expectations of the participants with respect to the language(s) of the input materials to be processed.

For the last three functions, the BIA model assumed that (by inhibiting non-target language words) the language nodes were able to facilitate the selection of words from the target language after language non-selective access took place, but could not enforce language selective lexical access from the very beginning of word recognition. Indeed, simulations show that even with strong pre-activation of the language nodes and strong top-down effects to the word level, word candidates from the surpressed language can often be recognized.

Before we discuss a number of limitations and problems of the BIA model, we will first summarize six types of empirical findings that the BIA model has been able to account for. We will restrict ourselves to the bilingual domain (also see Dijkstra et al., 1998(a); for additional monolingual simulations, see Van Heuven, 2000; Van Heuven, Dijkstra, Grainger and Schriefers, 2001).
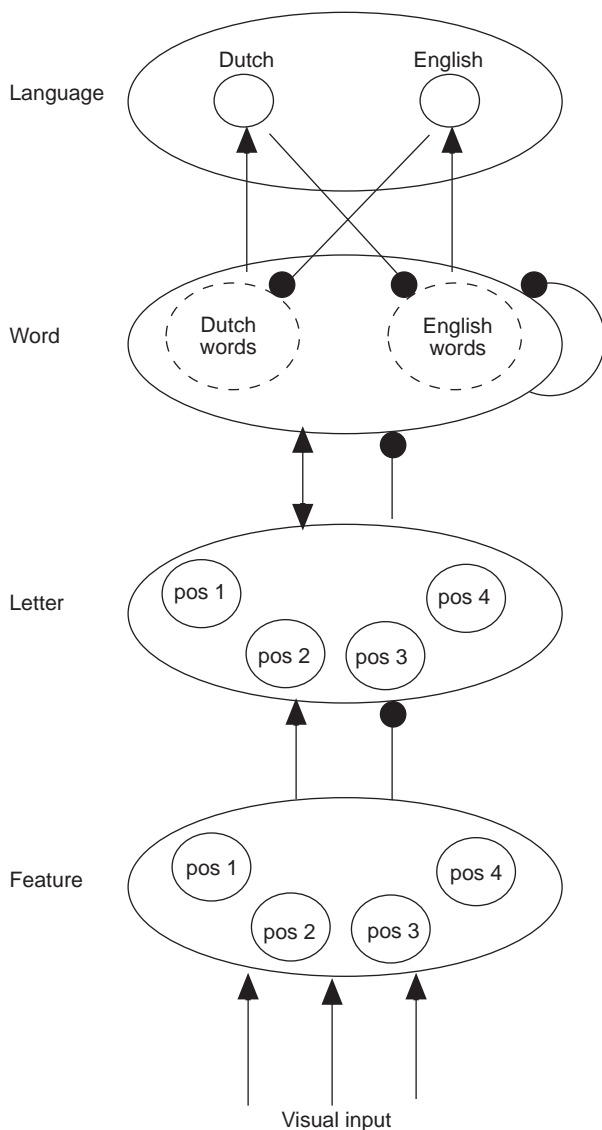

Figure 1. The Bilingual Interactive Activation (BIA) model for bilingual word recognition. Arrowheads indicate excitatory connections; black filled circles indicate inhibitory connections.

nodes in the other language. The language nodes collect activation from words in the language they represent and inhibit active words of the other language. The activation of the language nodes reflects the amount of activity in each lexicon.

In the BIA model, the language nodes fulfill four main functions (Dijkstra and Van Heuven, 1998). In the first two functions, language nodes serve as *linguistic representations*, in the last two as *non-linguistic functional mechanisms*. First, the language nodes serve as *language tags*, or language labels, representations that indicate the language to which an item belongs. These representations implement the observation that mature language users know to

---

[1] This account implies that the speed of activation of each language node depends on the number of items connected to that node (thus on vocabulary size in the concerned language), and on the (subjective) frequency of the word candidates involved. Such a mechanism is likely to fail if it is applied to beginning bilinguals without making additional assumptions (consider, for instance, what happens if only one L2 word has been acquired).

2.1 Simulating neighborhood density effects within and between languages

In neighborhood studies, it is investigated how the recognition of a target word is affected by the number of words that are similar to it. These similar words are called "neighbors". An orthographic neighbor is any word differing by a single letter from the target word respecting length and letter position (Coltheart et al., 1977). For instance, CORK and WORD are both neighbors of WORK. Monolingual word identification and naming turn out to be sensitive to the number (density) of neighbors and their frequency (Andrews, 1989; Grainger, O'Regan, Jacobs and Segui, 1989; Grainger and Segui, 1990; Snodgrass and Mintzer, 1993). In bilingual studies, effects of the number of orthographic neighbors were used as indexes of the relative influence of non-target language words on target word recognition in different experimental tasks and conditions. Target words themselves belonged only to one language (i.e., there were no interlingual homographs, homophones, or cognates in the stimulus list).

Because neighborhood density effects are assumed to arise during word identification, cross-language manipulation of neighbors allows a test of the independent versus integrated lexicon hypothesis. According to the integrated lexicon hypothesis, recognition of a target word will be affected by both target and non-target language neighbors, at least when lexical access is language non-selective. According to an independent lexicon hypothesis, recognition of the target word should not be affected by inter-lexical neighborhood density, because there are no direct interactions between the two lexicons. Thus, only a non-selective-access integrated-lexicon model such as BIA predicts that target word recognition is influenced by orthographic neighbors from both languages. Other models predict no effects of non-target language neighborhood density on target word recognition. As such, studies involving neighborhood manipulations offer a way to test the structural and processing issues discussed above at the same time.

In a series of progressive demasking and lexical decision experiments involving Dutch-English bilinguals, Van Heuven et al. (1998) examined neighborhood interference effects on target word recognition. They manipulated the number of orthographic neighbors of target words in the same and the other language of the bilinguals. Cross-linguistically, increasing the number of Dutch orthographic neighbors systematically slowed down response times (RTs) for English target words. Within the target language itself, an increase in

neighbors consistently produced inhibitory effects for Dutch target words and facilitatory effects for English target words. Monolingual English readers also showed facilitation due to English neighbors, but no effects of Dutch neighbors. Why opposite effects were observed of English (facilitation) and Dutch (inhibition) neighbors was not completely clear, but may have been a consequence of differences in the lexical organization of English and Dutch. Whatever explanation is favored, the results of the study indicate that during the presentation of a target word, neighbors from both languages are activated that all affect target recognition. This is evidence that the lexicon of bilinguals is integrated and non-selective in nature, at least with respect to orthographic codes.

Van Heuven et al. (1998) and Dijkstra et al. (1998(a)) showed that the empirically obtained neighborhood density effects could be simulated by the BIA model. A language non-selective access variant of the BIA model (without cross-linguistic top-down inhibition from the language nodes to the word level) already provided a good fit to the data of four different experiments when the English frequency range was restricted (implicitly assuming a lower subjective frequency for English L2 words). However, the introduction of asymmetric cross-linguistic top-down inhibition led to an even better fit across the four experiments. This provides a demonstration of the third function of the language nodes mentioned above, namely that they help to account for differences across experiments. Interestingly enough, the BIA model was able to account for the opposite effects of English and Dutch neighbors by using a combination of reduced English frequencies (relative to Dutch) and asymmetric top-down inhibition.

2.2 Simulating shifting neighborhood effects across an experiment

In one of their neighborhood density experiments, involving the progressive demasking technique, Van Heuven et al. (1998) observed that the pattern of RTs in their neighborhood conditions (orthogonally manipulating the number of neighbors in Dutch and English) changed for high-proficiency participants across the four parts of the experiment. In this blocked experiment, some participants saw the Dutch items first, while others saw them only after the English items. An inhibitory effect of English neighborhood size on Dutch target identification was observed, which decreased when the English items had been presented earlier. Thus, the Dutch RTs were strongly inhibited by the number of English

neighbors when the Dutch words were presented in the first block. The effect was especially strong in the first half of the first block, but was reduced in the second half. For the group of participants who saw the Dutch words only after the English words, the effect of English neighbors was reduced even further and turned into facilitation in the last half of the second block. This finding suggested that high-proficiency participants in some way modulated the relative contribution of the two languages. In the beginning of the experiment, participants expected their English proficiency to be tested and therefore kept their English lexicon in an active state (Grosjean, 1997). After being tested with English words, they could focus on the Dutch task and the relative activation of English may have dropped.

This account was supported by simulations with the BIA model in which only the settings were varied of the parameter that controlled the amount of inhibition exerted by the language nodes on the word level. Increasing the amount of asymmetric top-down inhibition across time from zero to larger values, the pattern of results could be fitted quite well. This can be considered as a demonstration of the fourth function of the language nodes discussed above.

2.3 Simulating masked priming effects in bilinguals

The BIA model has also been able to simulate the effects of masked orthographic priming within and between languages observed in a study by Bijeljac-Babic, Biardeau and Grainger (1997). In this study, French-English bilinguals made lexical decisions to L1 or L2 target words preceded by masked orthographic prime words from the same or a different language. In their first experiment, the target language was English, in the second experiment it was French. The target words in both experiments were low-frequency words or nonwords that were orthographically legal and pronounceable. The prime words were high-frequency words (presented for 57 ms) that either belonged to the same language as the target words or to the other language. In both experiments it was found that when prime and language were from the same language, target word decision times were inhibited for orthographically related primes relative to orthographically dissimilar primes. For example, when the target language was English, the prime-target combination "real – HEAL" led to slower RTs than the combination "roof – HEAL". This finding suggests that any facilitation effects that might arise due to form priming (overlap in letters) were annihilated by inhibition effects due to lexical competition. More

relevant, however, is that when prime and target were words from different languages ("beau – BEAM"), inhibition effects were found as well. This indicates that lexical knowledge from the other language affected target recognition, which provides evidence supporting language non-selective access to the bilingual lexicon. The response patterns of high-proficient bilinguals showed a considerable within-language inhibitory priming effect (28 ms) and a trend towards even larger cross-language inhibition (43 ms).

To simulate these results, a combined French–English lexicon was incorporated in the BIA model (Dijkstra et al., 1998(a)). Subsequently, standard model simulations were run with a fixed recognition threshold. At the first processing cycle, the prime word was presented to the model and at the third cycle it was replaced by the target word. In correspondence with the empirical data, the model produced longer average RTs in the within-language prime condition and the trend towards larger inhibition effects in the between-language prime condition. Note that it was not necessary to assume top-down inhibition effects (asymmetric or otherwise) from the language nodes to the word level to produce these inhibition effects.

2.4 Simulating L2-proficiency differences in masked priming with bilinguals

In the second masked priming experiment of their study, Bijeljac-Babic et al. (1997) examined the data patterns for three different L2-proficiency groups. The pattern for more advanced bilinguals was the same as described above: clear inhibition effects in the related prime condition relative to the unrelated control condition. Beginning bilinguals, however, showed smaller RT differences between the related and unrelated priming conditions, while there were no priming effects in monolinguals. The BIA model was able to simulate the three patterns of results when it was assumed that the monolinguals knew a few words of the foreign language (English) after all (Dijkstra et al., 1998(a)). If the monolinguals really did not have any foreign knowledge, the L2 words should be treated as nonwords and according to the BIA model, this should have produced a facilitation effect in the related condition relative to the unrelated condition (due to form overlap), rather than a null-effect. (Such masked priming facilitation effects for monolinguals and beginning bilinguals have indeed been observed in other recent work in our laboratory). Note that again simulations were simple; no top-down inhibition was assumed.

2.5 Simulating effects for interlingual homographs in a go/no-go task

Many bilingual word recognition studies investigating the issue of (non-)selective access have used interlingual homographs or cognates as stimulus materials. Interlingual homographs in a strict definition are words that are identical across languages with respect to their orthography but not their meaning (nor, most often, their phonology). Other terms used are "interlexical homographs" or "false friends". An example is a word like ROOM, which means "cream" in Dutch. Cognates are defined here as words from two languages that are identical in orthographic form and largely overlap in meaning, such as FILM in Dutch and English. The argumentation for using interlingual homographs and cognates in bilingual research has been as follows. Consider the Dutch-English interlingual homograph LIST (meaning "trick" in Dutch). If the recognition of the English reading of the homograph by Dutch-English bilinguals is affected by the Dutch reading, then response latencies should be different from those to one-language control items that are matched to the interlingual homograph in frequency, length and other characteristics (e.g., MILK). For instance, the requirement of selecting either the English or the Dutch reading of a homograph might induce inhibitory effects relative to a control. However, if recognition proceeds in a language selective way, then no RT differences between homographs and controls would be expected, because the Dutch reading of the interlingual homograph would not be activated at all and would not affect lexical selection based on the English lexicon.

To investigate this issue, Dijkstra, Timmermans and Schriefers (2000(b)) made use of a go/no-go paradigm. In this paradigm, participants had to react only when a presented word belonged to a prespecified target language. Participants reacted only when they identified either an English word (English go/no-go) or a Dutch word (Dutch go/no-go), but they did not respond if a word of the non-target language (Dutch or English, respectively) was presented. The list of presented words consisted of a mixture of interlingual homographs and one- language Dutch and English control groups. The participants were informed that some of the words in the list could be both English and Dutch: interlingual homographs. Such homographs could belong to three types: high-frequent in English (HE) and low-frequent in Dutch (LD); low-frequent in English (LE) and high-frequent in Dutch (HD); and low-frequent in both languages. Examples of words in the three groups are LIST (HE/LD), BRAND (LE/HD) and GIST (LE/LD).

In both go/no-go tasks, clear inhibition effects arose for homographs relative to one-language controls. Even in the Dutch go/no-go task for Dutch-English bilinguals performing in their native language, participants were unable to completely exclude effects from the non-target language on homograph identification. Target-language homographs were often "overlooked", especially if the frequency of their other-language competitor was high. In the Dutch go/no-go task, participants did not respond to low-frequency items belonging to their native language in about 25 percent of the cases! Inspection of cumulative distributions showed that if they did not respond after about 1500–1600 ms, they did not respond within the time window of two seconds anymore.

To simulate these results with the BIA model, it was assumed that a homograph is represented in the mental lexicon twice, once for each language (Dijkstra and Van Heuven, 1998). Each representation has a resting level activation depending on its frequency of occurrence in the language that it belongs to. Furthermore, these representations compete with words of both languages in a standard fashion. Using asymmetric top-down inhibition from the Dutch language node to English words, the pattern of correct responses (in percentages) for homographs and control words in the three frequency conditions of the Dutch go task was reproduced by the BIA model reasonably well. The percentage of correct responses was clearly dependent on the relative frequency of both readings of the interlingual homographs. However, simulating the response times for the different types of homographs turned out to be more difficult.

2.6 Simulating language of previous item effects

Von Studnitz and Green (1997) had German-English bilinguals perform a generalized lexical decision task, in which they gave a "yes" response if a presented item belonged to German or English, and a "no" response if it was a nonword. Words in the list were sometimes preceded by other words from the same language, and sometimes by words from the other language. Von Studnitz and Green observed that the bilinguals' RTs on switch trials were a significant 17 ms slower than on non-switch trials. Significant language switching effects of considerable size were also found in a reanalysis of the generalized lexical decision data of the Van Heuven et al. (1998) neighborhood study as well. Dutch words preceded by nonwords (571 ms) were responded to faster than those preceded by English words (586 ms) and slower than those preceded by Dutch words (556 ms).

English words preceded by nonwords (621 ms) were responded to faster than those preceded by Dutch words (641 ms) and slower than those preceded by English words (606 ms).

In a psychophysical study with themselves as the two participants, Grainger and O'Regan (1992) made lexical decisions to repeatedly presented words of four letters and nonwords that were preceded by unrelated French and English prime words. At presentation durations for which the primes could be identified, language prime effects arose that could be accounted for by assuming that primes from a different language interfered with target processing.

Such findings can, in principle, be explained in terms of the BIA model by assuming that, due to the interaction of the language nodes with the orthographic lexical representations, target word recognition is affected differently by preceding words from the same and from another language (Dijkstra, Van Jaarsveld and Ten Brinke, 1998(b), p. 63–64). For instance, on trial t an English word would activate the English language node, and on trial t + 1, this language node would feed activation back to all English words and/or inhibit all words from the Dutch lexicon. This hypothesis could hold for both types of studies just reviewed, including item pairs across trials ("previous item effects") or within-trial prime-target pairs (e.g., in unmasked or masked priming). However, because evidence from several recent studies (Von Studnitz and Green, 1997, 2002; Thomas and Allport, 2000) indicates that the major source of switch costs must be positioned outside the mental lexicon, this account is probably not correct (see also section 4.8).

3 Limitations and problems of the BIA model

Thus, in spite of the successful simulations of the BIA model described in the previous section, empirical studies suggest that some of the proposed mechanisms with respect to consecutive item effects and language switching may be wrong. In addition, there are many aspects of bilingual word recognition that are not fully accounted for by the model. To mention just a few of these:
- there are no phonological or semantic representations in the model;
- the representation of interlingual homographs and cognates is underspecified;
- representational and functional aspects with respect to the language nodes are confounded;
- there is only a very limited account of how non-linguistic and linguistic contexts affect bilingual word recognition;
- there is no detailed description of how participants

perform a particular task, for instance lexical decision;
- the relationship between word identification and task demands is underspecified.

The majority of points in this list call for the extension of the BIA model with additional representations and processing components. In addition, the model must be differentiated and adapted with respect to the language node representations. In the following sections, we will describe the revised model, which we will call the BIA+ model. To the extent that the BIA model is "nested" in the BIA+ model, the earlier simulations on the basis of orthographic representations only are still valid for the BIA+ model. However, as we shall see, implementing lexical phonological and semantic representations computationally poses serious problems, and such implementation problems are even more serious for the task/decision component (in fact, at present there is no available monolingual model that has incorporated both aspects in a satisfactory way). We will, therefore, at this moment, follow Green (1998) and Grosjean (1997) in providing only a verbal analysis for these components of the model. We also acknowledge that there are many other aspects of bilingual word recognition that should be considered in the future (for instance, how the model would develop over time and during learning), because they might affect theoretical perspective and focus.

4 The BIA+ model for bilingual word recognition

The BIA+ model is graphically represented in Figure 2. It incorporates the BIA model as a special case, except that the function of the language nodes is adapted. The BIA+ model is also strongly affected by Green's (1998) ideas on task schemas and task control. This is evident in the distinction between a word identification system and a task/decision system. Green's (1998) Inhibitory Control (IC) model specifies the control that bilinguals have over the processing in their lexico-semantic system in different task conditions. Both Green (1998) and Dijkstra (1998) have pointed out that the relationship between the IC model and the BIA model is rather complementary, with a larger focus on the task schema and bilingual language production in the IC model and on the bilingual lexico-semantic system and comprehension in the BIA model.

In separate sections, we will consider the BIA+ model with respect to the following points: (1) Representation and processing of orthographic, phonological and semantic codes; (2) Representation of interlingual homographs and cognates; (3) Language nodes; (4) Linguistic context effects; (5) Non-lin-
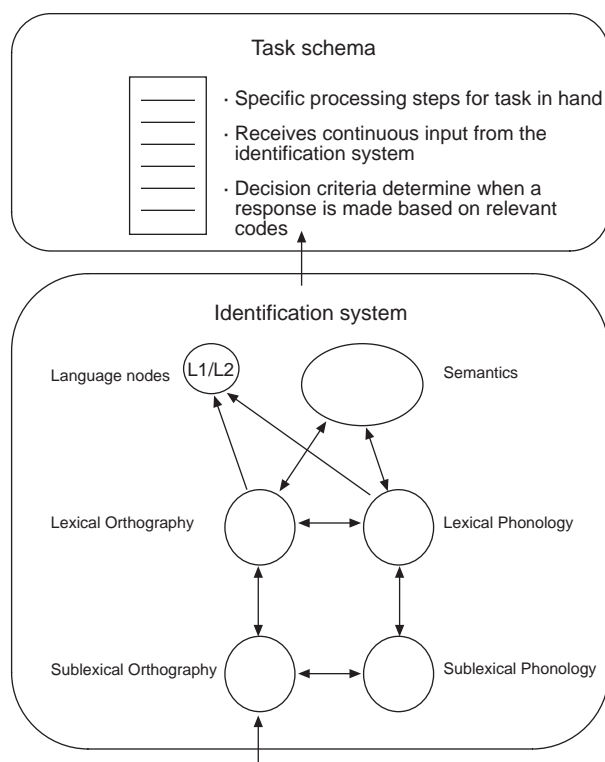
Figure 2. The BIA+ model for bilingual word recognition. Arrows indicate activation flows between representational pools. Inhibitory connections within pools are omitted. Language nodes could instead be attached to lemma representations between word form and meaning representations. Non-linguistic context only affects the task schema level.

guistic context effects; (6) Relationship between word identification and task demands; (7) Stimulus-response binding in lexical decision; (8) Stimulus-response binding in language switching.

4.1 Representation and processing of orthographic, phonological and semantic codes

Just like the BIA model, the BIA+ model proposes that the bilingual lexicon is integrated across languages and is accessed in a language non-selective way. However, the BIA+ model extends these assumptions from orthographic representations to phonological and semantic representations. In other words, bilingual word recognition is affected not only by cross-linguistic orthographic similarity effects, but also by cross-linguistic phonological and semantic overlap. This is in line with several recent studies demonstrating profoundly language non-selective effects with respect to all three codes (e.g., Dijkstra, Grainger and Van Heuven, 1999; Jared and Kroll, 2001).

Implementing phonological and semantic representations in an interactive activation model, however, poses serious problems for the modeler. For instance, in English the mapping between letters (graphemes) and phonemes is context sensitive, implying that a sublexical representation in terms of such units alone will not suffice. This is a well-known problem in the monolingual domain, for which several tentative solutions have been proposed (Dijkstra, in press). In a recent interactive activation model for monolingual word recognition, we included a new representational layer of orthographic and phonological representations as a solution to this problem, using the so-called Onset–Nucleus–Coda (ONC) scheme for both code types (Plaut, Mc-Clelland, Seidenberg and Patterson, 1996; Van Heuven, 2000). According to this scheme, a word like STRAND, for instance, can be represented by an onset cluster STR, a vowel nucleus A, and a coda ND. Assuming this type of representation for both orthography and phonology simplifies the implementation of a context sensitive mapping between the two types of codes. Preliminary simulation work indicates that this implementation type is successful in accounting for phonological effects (e.g., consistency effects) in monolingual word recognition, but more work is necessary to assess the feasibility of this approach for implementing phonological codes in the BIA+ model.

Activation of orthographic codes

When an input letter string is presented to the BIA+ model, the first stages of word recognition proceed in the same fashion as in the BIA model. A number of lexical orthographic candidates are activated in parallel depending on their similarity to the input string, and on the resting level activation of the individual items (and therefore dependent on subjective frequency, recency of use, L2 proficiency, etc.). Because L2 representations are on average of a lower subjective frequency than L1 codes, they are activated somewhat more slowly than L1 representations. Next, activated orthographic word candidates activate their corresponding phonological, semantic and other (e.g., articulatory) representations.

According to the model it is the similarity of the input word to the internal lexical representations that determines their activation, not the word's language membership. The larger the overlap between the input string and a representation in the mental lexicon, the more the internal representation is activated. As a consequence, in the case of two languages with alphabetical writing systems, the number of activated orthographic candidates is determined by factors such as the neighborhood density and fre-

quency of the target word and its within- and between-language neighbors. If the two languages differ with respect to their input codes (e.g., letter sets), the activated set of neighbors may become much smaller. For instance, assume that a presented item contains language specific bigrams or diacritical markers. Due to their presence, the initial set of lexical candidates that is activated may become restricted to one language (cf. Mathey and Zagar, 2000, p. 200). This is a bottom-up effect. We deem it unlikely that for words that are well-known by the bilingual this kind of information is often used in a top-down way (e.g., the bilingual noticing a particular language specific bigram and using it for language decision), because the bottom-up recognition process will usually be much faster. (Of course, if all German words in a language decision task begin with a capital letter, then participants can redefine the task of responding to words as one in which they respond to capital letters.)

Note that it follows logically that no "orthographically similar" word candidates can be activated across language pairs that do not share orthography at all (e.g., Chinese and English), even though effects of phonological similarity might still occur for such language pairs. In other words, when particular input aspects are language specific, we will (of course) find evidence of language specific access (e.g., Chinese orthography will not induce much Latin letter activation).

For interlingual homographs and cognates with different orthographic forms (and therefore two representations) across languages, the degree of code activation of the non-target reading also depends on the degree of cross-linguistic code overlap (see section 4.2). This raises the question of how recognition of the Dutch word TOMAAT by Dutch-English bilinguals is affected by its similarity but non-identity to the English word TOMATO. Van Hell and Dijkstra (in press) had trilinguals with Dutch as their L1, English as their L2 and French as their L3 perform a word association task or a lexical decision task in their L1. Stimulus words were (mostly) non-identical cognates such as TOMAAT or noncognates. Shorter association and lexical decision times were observed for Dutch-English cognates than for noncognates. For trilinguals with a higher proficiency in French, lexical decision responses were faster for both Dutch-English and Dutch-French cognates. Thus, even when their orthographic and phonological overlap across languages is incomplete, cognates may be recognized faster than noncognates.

For French-Spanish bilinguals, Font (2001) has found that in lexical decision cognates differing in one letter between languages (called by her "neighbor cognates") are still facilitated but significantly less so than identical cognates. Furthermore, the amount of facilitation observed depended on the position of the deviating letter in the word. Neighbor cognates differing at the end of the word (e.g., French TEXTE – Spanish TEXTO) were facilitated more than neighbor cognates with the different letter inside (e.g., French USUEL – Spanish USUAL). In fact, facilitatory effects for the latter type of cognate disappeared and effects tended towards inhibition when such cognates were of low frequency in both languages. Similar patterns of results were found in L1 and L2 processing. These results suggest that the size of cognate and interlingual homograph effects depends on their degree of cross-linguistic overlap (cf. Cristoffanini, Kirshner and Milech, 1986), just like the BIA+ model predicts. However, in order to account for position-specific effects of mismatches, the letter to word connections in the BIA+ model must be differentiated with respect to position, or a new letter coding scheme must be introduced (also see Dijkstra, in press).

Activation of phonological and semantic codes

When (sublexical and lexical) orthographic representations become active, they start to activate associated phonological and semantic representations. As a consequence, such phonological and semantic representations of both languages are activated slightly later during word reading than the orthographic representations of these languages (Ferrand and Grainger, 1993). Because their activation will depend, among other factors, on subjective frequency, this implies that L2 phonological and semantic codes will be delayed in activation relative to L1 codes. We will call this the "temporal delay assumption". However, the different codes will often interact (resonate) over time and their respective identification will take place close in time. As a consequence of time course differences, the relative contribution of codes to the decision may be modulated by adapting temporal deadlines rather than by refiguring the task schema. For instance, when pseudohomophones are present in a monolingual stimulus list, phonological effects could be attenuated or excluded by speeding up RTs and basing the response on the orthographic representation only. However, mutual activation of different codes may make phonological effects difficult to eliminate.

Two consequences of the temporal delay assumption are that (1) cross-linguistic effects will generally be larger from L1 to L2 than in the opposite direction; (2) an absence of cross-linguistic phonological and semantic effects for different words could occur

if task demands allow responding to faster codes (for instance, orthographic L1 codes), giving slower codes no chance to affect the response times.[2] The systematic task dependence of cross-linguistic effects is shown clearly in a number of recent empirical studies (e.g., Dijkstra et al., 1998(b); Dijkstra, Grainger and Van Heuven, 1999; Lemhöfer and Dijkstra, submitted).

We will illustrate this point by comparing the study by Dijkstra et al. (1999) with that by Lemhöfer and Dijkstra (submitted). Dijkstra, et al. (1999) examined the effects of different types of code overlap in interlingual homographs and cognate processing. Dutch-English bilinguals performed an *English lexical decision* task with English words varying in their degree of orthographic (O), phonological (P), and semantic (S) overlap with Dutch words. Their six different test conditions are exemplified by the following items: "hotel" (overlap in S, O and P codes), "type" (SO), "news" (SP), "step" (OP), "stage" (O) and "note" (P). The first two conditions (SOP and SO conditions) may be called "cognates", while the last three conditions contain "interlingual homographs" (OP and O conditions) and "interlingual homophones" (P condition). Lexical decisions were facilitated by cross-linguistic orthographic and semantic similarity relative to control words that belonged only to English. In contrast, phonological overlap produced inhibitory effects. A very similar pattern of results was found using a different task (progressive demasking), but no systematic differences between test and control conditions arose for American English monolinguals.

We conclude that in this English lexical decision task, where L2 (English) was the target language, cross-linguistic effects arose for L1-L2 (Dutch-English) homographs with respect to all three types of representations. Because English was the target language in this task, a proper or safe task execution implied verifying the English language membership of possible word candidates, even when Dutch codes would have become available faster than English ones. In other words, Dutch codes had time to establish themselves and exert effects on later available English codes that were necessary for responding.

Lemhöfer and Dijkstra (submitted) presented the same stimulus materials to Dutch-English bilinguals in a *generalized lexical decision* task. In this task, participants responded with "yes" to both English and Dutch words, but with "no" to nonwords. In contrast to English lexical decision, participants in this task can use both Dutch and English lexical representations as a reliable basis for responding. Thus, in this task, cross-linguistic effects will arise only to the extent that L1 and L2 codes can affect each other before the fastest codes (usually Dutch ones, according to the temporal delay assumption) are retrieved and responded to. The results of this study were straightforward: no facilitation effects arose for interlingual homographs relative to controls, while cognates were facilitated. The pattern of results for homographs indicates that responses were based upon the fastest available code, usually the Dutch orthographic code, while cross-linguistic overlap with respect to semantics in the case of cognates apparently can be used to speed up the response. This implies a special type of representation for cognates, possibly with a strong feedback connection from semantics to orthography (cf. Pecher, 2001; Reimer, Brown and Lorsbach, 2001).

The result patterns of the two studies are accounted for in the BIA+ model by assuming that the task differences can lead to a different use of lexical codes while the activation pattern in the word identification system itself is not changed. This assumption is more parsimonious (and easier to test) than the assumption of changes in both identification and task/decision systems (an assumption that is made by the BIA model).

4.2 Representation of interlingual homographs and cognates

Because the BIA model included only orthographic representations, simulations involving orthographically identical interlingual homographs could only be conducted if two representations were assumed for such items (see section 2.5). This choice, however, was practically motivated and not based on solid empirical evidence. Note that multiple representations for form-identical items are not necessarily incompatible with an integrated lexicon account, even though they fit more naturally within a bilingual word recognition account that assumes separate lexicons for each language. In this section, we will investigate this issue in some detail. As an example, do ROOM (a Dutch-English interlingual homograph) and FILM (a Dutch-English cognate) have shared or distinct representations in English and Dutch?

A comparison of recent studies suggests that orthographically identical interlingual homographs have distinct orthographic representations in each language (Dijkstra et al., 1998(b); Dijkstra et al.,

---

[2] There are several other ways in which these null-effects may have arisen. They may, for instance, be due to participant strategies or the effects of different codes (e.g., orthography and phonology) may cancel each other.

1999; Lemhöfer and Dijkstra, submitted). In the previous section, we discussed the study by Dijkstra et al. (1999), who found cross-linguistic effects for L1-L2 (Dutch-English) homographs with respect to orthographic, phonological and semantic codes in a bilingual lexical decision task where L2 (English) was the target language. Dijkstra et al. (1998(b)) manipulated a different aspect of interlingual homographs, the relative frequency of the two readings of the homographs. In their second experiment, Dutch-English bilinguals also performed an English lexical decision on interlingual homographs, exclusively English control items, and nonwords. The experimental stimuli also included Dutch words requiring a "no" response (i.e., the Dutch words had to be treated as "English nonwords"). Relative to English control items, strong inhibitory effects were obtained for interlingual homographs that were dependent on the relative frequency difference of the two readings of the homograph. The inhibitory effect was large when the Dutch reading of the homographs had a high frequency relative to the English reading.

The results of these two studies are paradoxical. On the one hand, the observed cross-linguistic facilitatory effects for orthographic identity and semantic overlap in interlingual homographs suggest that such representations are shared for items occurring in two languages. Shared representations for interlingual homographs within an integrated lexicon fit well within an interactive activation account of bilingual word recognition. On the other hand, the observed differential contributions of the Dutch and English frequencies of interlingual homographs in the second study suggest that these items have different lexical representations, each characterized by its frequency in the language it belongs to.

Is there any evidence that allows us to distinguish views that assume one (orthographic) representation or two (orthographic) representations for interlingual homographs? If there is only one orthographic representation, (orthographic) frequency effects should be cumulative under circumstances in which both readings of the homograph can be used for responding. This should be the case then in a generalized lexical decision task, such as that of Experiment 3 in Dijkstra et al. (1998(b)). In this task, participants can respond to the fastest identified reading of the interlingual homograph and do not need to make a distinction in order to perform the task. The one-representation hypothesis predicts, first, that the responses to interlingual homographs should be faster than the fastest of the control conditions. Second, the cumulative frequency effect should be especially visible if, keeping the frequency of one reading of the homograph constant, the frequency of

the other reading is varied. The obtained result pattern, however, was not in line with this view. First of all, the facilitation effect for the HFE-LFD condition is only 9 ms relative to the English control condition, and for the HFE-HFD condition just 12 ms relative to the Dutch control condition. Furthermore, for the LFE-HFD condition a 22 ms inhibition effect arose relative to the fast Dutch control condition, and for the LFE-LFD condition even 28 ms of inhibition arose. The relatively small facilitation in the first two conditions could be accounted for by assuming that the cumulative frequency effects are only limited because log frequency rather than frequency is important; however, the inhibitory results for the low-frequency conditions cannot be explained in such a way.

The available results are easier to account for in the two-representations view because, according to this view, responses to interlingual homographs cannot only be slower than to controls, but to a limited extent also faster as an effect of "statistical facilitation" (Raab, 1962). Statistical facilitation can arise if the response time is determined by the first available reading of the homograph. In most trials this would be that of the high-frequency reading of the homograph, but sometimes the low-frequency reading might win the race, making the resulting mean RT just somewhat faster than that of the high-frequency control (also see Lemhöfer and Dijkstra, submitted).

What can be said about the representation of cognates, i.e. interlingual homographs that have both orthographic and meaning overlap across language? The available studies suggest that cognates have a special representation. Apart from the studies by Van Hell and Dijkstra (in press), Dijkstra et al. (1999) and Lemhöfer and Dijkstra (submitted) that we have already discussed, there are many earlier studies supporting this suggestion. For instance, Kirsner and colleagues (e.g., Lalor and Kirsner, 2000), and Sánchez-Casas and colleagues (e.g., Sánchez-Casas, Davis and García-Albea, 1992) have proposed that cognates may have a common morphemic representation across languages. However, more detailed information with respect to the relationship between their orthographic, phonological and semantic codes is necessary before a complete cognate representation can be implemented.

To conclude, the presently available evidence favors the view that interlingual homographs are represented by two (possibly partially overlapping) representations rather than one. Thus, it seems reasonable to continue with BIA+ simulations assuming two representations for interlingual homographs rather than one. Additional evidence about

cognate representation is necessary in order to implement these in models like BIA+.

4.3 Language nodes

In section 2, we discussed the various functions of the language nodes in the BIA model. We made two distinctions: between linguistic and non-linguistic functions of the language nodes, and between bottom-up (word to language) and top-down (language to word) components of these functions. In an earlier paper on the BIA model (Dijkstra and Van Heuven, 1998), we already indicated that combining all these aspects in one mechanism was probably too ambitious. Since that time, evidence has accrued suggesting that the linguistic and non-linguistic functions of the language nodes must be assigned to different levels of processing and that they may operate in a different way. More specifically, we propose to restrict the language nodes' functions to language membership representations within the identification system of which the activation level also reflects global lexical activity (because all activated words of one language feed activation forward to a language node). Being just representations, the language nodes can no longer function as language filters dependent on experimental factors nor as collectors of non-linguistic contextual pre-activation.

Differences between result patterns across experiments and non-linguistic context effects now need to be accounted for at a different level. We propose that these effects arise at the level of task specification and parameter settings for the decision in a task (see Figure 2). In other words, we propose a distinction between a word identification system and a task/decision system, analogous to Green (1998). Cross-experimental differences, cross-trial differences or participant expectations may be handled by a decision mechanism affecting the output of the word identification system before the response is made.

This proposal changes the nature of the language nodes, but they remain present as a useful construct in the model. Language users know to which language a word belongs, so there must be some sort of language tag or language membership representation. In language decision, participants are assumed to retrieve this representation to decide whether they should push one button or another. They cannot do this just by identifying the item, because even though it may very clearly belong to one language only, it is not clear which language that is before the tag has been retrieved. The language membership information could be retrieved via the item's lexeme (orthographic or phonological) or lemma (more abstract syntactic/semantic) representation. For the

sake of simplicity, the model currently does not incorporate lemma representations, implying that the language nodes are directly connected to lexical form representations. According to this viewpoint, the relative activation of the language nodes (some would say "of the languages") is completely dependent upon activation arriving from other linguistic representations, e.g., current lexical input and previous sentence context.

Studies indicate that language information becomes available rather late during (isolated) bilingual visual word recognition, usually too late to affect the word selection process. An example is the study by Dijkstra et al. (2000(b)), discussed already in section 2.5. Target-language homographs were often "overlooked" in a language go/no-go task if the frequency of their other-language component was high. A flattening of the cumulative RT distribution towards an asymptotic value in these experiments suggests that recognition of the homograph reading from the non-target language in some way "prohibited" the subsequent recognition of the target language reading (e.g., after recognition, all other lexical candidates may be suppressed). Thus, selection of one reading of the interlingual homographs took place rather late during processing. Clearly, the system must at some time select only one lexical item, but apparently the language of that item could not aid that selection. More likely, determination of the item's language depended on lexical selection having taken place. In addition, it did not seem possible to focus on the target reading only and discard the non-target language reading. One reason for this may be a tendency that the word rather than its language label triggers the response (see also section 4.7).

In section 4.5 we will discuss some other studies on word reading that suggest that the influence of language membership on item recognition is relatively small, implying it is generally available too late to affect word identification. If this observation proves to be true in further research, it means that the feedback parameter from language nodes to the word level must be set rather low. One may wonder how this can come about in an interactive activation model such as BIA+. A possibility (borrowed from the monolingual language production literature) lies in the nature of the mapping between words and language membership. Each word is connected to one language node, but each language node is connected to thousands of words. Therefore, if the amount of activation that can be distributed between units at different levels is a constant, the feedback activation from language node to word level will be much smaller per word unit than the feedforward activation from word level to language node.

4.4 Linguistic context effects

The BIA+ model makes a distinction between a word identification system incorporating linguistic representations, and a task/decision system incorporating non-linguistic task schema specifications (see Figure 2). In line with this distinction, it is assumed that linguistic context is able in principle to directly affect the word identification system, while non-linguistic context affects (settings in) the task/decision system. Linguistic context effects are defined here as the effects arising from lexical, syntactic or semantic sources (e.g., sentence context), whereas non-linguistic context effects are those effects that arise from instruction, task demands or participant expectancies.[3] The automatically operating word identification system is part of a much larger Language User system in which sentence parsing and language production also have a place (Dijkstra and Kempen, 1984; Levelt, Roelofs and Meyer, 1999). We propose that the word identification system is interactive not only with respect to different codes but also relative to the sentence parsing system. As a consequence, BIA+ predicts that the recognition of words in sentence context is sensitive to syntactic and semantic context information from different languages, in a way that is analogous to monolingual word recognition in sentence context. In fact, such linguistic context information may exert serious constraints on the degree of language selective access that may be observed.

At present, few studies have addressed the bilinguals' recognition of words in sentence context (Altarriba, Kroll, Sholl and Rayner, 1996; Li, 1996; Elston-Güttler and Williams, submitted), but their results are in line with this viewpoint. We will discuss two of these here. In their first experiment, Altarriba et al. monitored the eye movements of Spanish-English bilinguals while they were reading English (L2) sentences that contained either an English (L2) or a Spanish (L1) target word. Sentences provided either high or low semantic constraints on the target words. An example sentence of the high constraint and Spanish target condition is "He wanted to deposit all his *dinero* at the credit union", where *dinero* is Spanish for "money". The experiment led to an interaction between the frequency of the target word and degree of sentence constraint for Spanish target words with respect to the first fixation duration, but not for English target words. Thus, when the Spanish target words were of high frequency and appeared in highly constrained sentences, the partici-

pants apparently experienced interference. This result suggests that sentence constraint influences not only the generation of semantic feature restrictions for upcoming words, but also that of lexical features. The high-frequency Spanish word matched the generated set of semantic features, but not the expected lexical features when the word appeared in the alternate language (Altarriba et al., 1996, p. 483). The same pattern of results was found in a second experiment, where the sentences were presented word by word using the rapid serial visual presentation (RSVP) technique and participants named the capitalized target word in each sentence.

The findings of this study are in line with the BIA+ model from at least two perspectives. First, linguistic sentence context was found to interact with target word recognition, as predicted by the model. Second, note that the observed data pattern showed an interaction of word frequency (a lexical information source) and the sentence constraint, and not of language membership and the sentence constraint. This suggests that (just like for isolated words) lexical characteristics are more important than language characteristics in the determination of word recognition in sentences.

Future studies should focus on disentangling such effects of lexical form features and language membership in sentence processing experiments. They should examine, for instance, to which extent the language itself of preceding words in the sentence can modulate the activation of target word candidates from a non-target language. Two viewpoints can be contrasted. One option is that the language nodes can be pre-activated by the sentence context and may function as a link between sentence and lexical levels. Via this link, the language of the preceding sentence context could affect the recognition of a target item independent of any additional syntactic or semantic effects. Another option is that such preactivation will not be effective, because language nodes cannot activate or suppress word activation to any considerable extent. This second option currently seems to be the most in line with the BIA+ model (and the study by Altarriba et al.), because it fits the argument put forward in section 4.3 that language information does not provide strong selection constraints on bilingual word recognition.

The second study we will discuss here is that by Elston-Güttler and Williams (submitted). It is of direct relevance to the BIA+ model, because they compared the bilingual processing of interlingual homographs in isolated word lists (see sections 2.5, 2.6 and 4.2) to that of the same stimulus materials incorporated in sentences. The authors first showed that, in word lists, visually-presented German-

---

[3] Effects of stimulus list composition could, in principle, derive from both sources.

English interlingual homographs primed lexical decisions on the English translation of their German reading (cf. De Moor, 1998; Van Heste, 1999). For example, the interlingual homograph "gift" primed the target word "poison". Next, the prime word was incorporated as the last word in a sentence like "The woman gave her friend an expensive gift". At a stimulus onset asynchrony (SOA) of 250 ms after the prime word, the target word "poison" was presented. In other conditions, the target item was a nonword. As before, the participants performed a lexical decision task on the target item. In contrast to the significant single word priming, no significant priming was obtained when the primes occurred in sentence contexts biased towards their English reading. As discussed above, such a result is not incompatible with the assumption of the BIA+ model that sentence context can interact with the bilingual word identification process.

In a second experiment, Elston-Güttler and Williams examined cases of L1-L2 relationship occurring at a deeper level than the form-level relations examined in Experiment 1. Using the same procedure as before, they presented German-English bilinguals with a sentence such as "The lawyer tried very hard to defend his client in court", followed by a target word like "meal". In this case, the prime "court" can be translated into German as "Gericht". The German word "Gericht" is ambiguous, and can also mean "meal". Interestingly enough, priming effects were now obtained in both the single-word and sentence conditions. The authors argued that this suggests that L1 representations that are activated by translation links are less resistant to sentence and language context constraints than those that depend on form-level similarity.

The results of this study are to a considerable extent compatible with those obtained in the monolingual domain. Studies by Swinney and others (e.g., Swinney, 1979; Tanenhaus, Leiman and Seidenberg, 1979) originally suggested that even in the presence of a strong biasing context, both readings of (intralingual) homographs were briefly activated. Later research has complicated this picture by showing that the more frequent (dominant) meaning of an ambiguous word may be more accessible than the other one (e.g., Duffy, Morris and Rayner, 1988), and that sentential context may make the non-dominant meaning as accessible as the dominant one (Lucas, 1999). Furthermore, under some conditions, only the contextually appropriate meaning appears to be activated (Tabossi and Zardon, 1993).

BIA+ proposes to use such monolingual results to formulate predictions for bilingual research based on the assumption that syntactic and semantic effects are language non-selective, just as word recognition is language non-selective. For instance, syntactic constraints will affect word recognition irrespective of the language of sentence context or target word, and syntactic lexical and phrasal categories may to some extent be language independent. Thus, effects of cross-linguistic syntactic priming between L1 and L2 are expected to occur.

Thus, the BIA+ model proposes that linguistic context, in particular syntactic and semantic context, may directly affect bilingual word activation via the word identification system. In the next section, we argue that non-linguistic information is more likely to affect the task/decision system. In fact, we will argue that word activation is not modulated by non-linguistic context.

4.5 Non-linguistic context effects

In principle, non-linguistic context (such as the participants' expectations on the basis of instruction or task demands) could affect a language non-selective word recognition system in several ways. One option is that both linguistic and non-linguistic context information can modulate the relative activation of items in the target and non-target language after the initial stages of lexical processing. For instance, context information might inhibit (i.e., reduce the activation of) lexical candidates or lemmas in the irrelevant language (cf. BIA model by Dijkstra et al., 1998(a),(b); IC model by Green, 1986, 1998) or induce a flexible modulation of the relative activation of lexical candidates in the two languages (Grosjean, 1997). A second option is that non-linguistic context information does not affect the activity in the identification system itself, but leads to an adaptation of decision criteria only. A third option is that adaptations in both activation levels and decision criteria are possible (as is the case in BIA and IC models).

The BIA+ model proposes that (in contrast to linguistic effects) non-linguistic context effects in reading can be accounted for only by the second option. Thus, effects that are explained by relatively late effects of top-down inhibition in the BIA model are accounted for by decision criteria adaptations in BIA+. So far, there seems to be little pertinent evidence in favor of top-down inhibition of active lexical candidates on the basis of task demands or participant strategies. In other words, it appears that the bilingual word identification system is encapsulated relative to decision level factors. Thus, for bilingual word recognition we reject the conclusion drawn by Gerard and Scarborough (1989, p. 314) that "bilinguals can exercise a considerable degree of control over access to lexical information".

How can the hypotheses of top-down inhibition or adaptation of decision criteria be contrasted? The answer depends in part on the specification of each view. Let us assume that top-down inhibition occurs in a particular experimental situation if the suppression of activation of items of one language leads to a better overall performance. Suppression here means that the activation of these items is reduced and that, as a consequence, their influence on other lexical candidates is minimized. We should then be able to observe the effects of such top-down inhibition by comparing two experimental situations that are as similar as possible except that one of them is optimized if unwanted items are suppressed, while the other does not benefit from (or is harmed by) such suppression.

A few recent studies seem to allow such a comparison, both within- and between-experiments. For instance, in the homograph studies by Dijkstra et al. (2000(b), Experiments 2 and 3) discussed earlier, participants would have given the fastest RTs to homographs and control items if they had suppressed items belonging to the non-target language (the same holds for Experiment 2 in Dijkstra et al., 1998(b)). Such an optimization by the participants would have led to RTs to homographs that were just as fast as to controls. Instead, strong inhibition effects were obtained indicating that suppression of the non-target language candidates was not possible. (The "inhibition effects" that were observed reflect competition of the two readings of the interlingual homographs rather than the suppression of one of them. Thus, they are not "top-down" inhibition effects as defined above.) Analogously, Dutch phonology was not relevant for the English lexical decision to be made in the experiments in Dijkstra et al. (1999). Nevertheless, it was activated and induced inhibitory effects on the response. In other words, parallel bottom-up activation of lexical candidates from both languages is so strong that suppressing it seems practically impossible.

A similar argument is found in the English lexical decision study by Dijkstra, De Bruijn, Schriefers and Ten Brinke (2000(a)). This experiment consisted of two parts. In part 1 of this study, which included only English words and nonwords, the RTs of Dutch-English bilinguals to interlingual homographs were just as fast as to English control items. In part 2, purely Dutch (non-target language) items were introduced (requiring a "no" response) that led to clear inhibition effects for interlingual homographs relative to controls. However, RTs to English control items were only marginally affected by the transition (581 ms to 592 ms, significant at p < .05 only in the item analysis), indicating little change in activation

going from part 1 to part 2. The observed small effect could be explained by a slight change in the decision process, as a consequence of the extra caution needed to exclude purely Dutch words.

Indeed, a comparison of the control conditions across experiments in Dijkstra et al. (1998(b)) and in De Groot, Delmaar and Lupker (2000), strikingly confirms this analysis. The three experiments in the study by Dijkstra et al. were conducted using different participants from one population and contained mostly the same test materials. Nevertheless, the experiments showed strong and systematic effects on interlingual homograph recognition of the inclusion of Dutch words (Experiment 2 vs. Experiment 1) and changes in task demands (Experiments 1 and 2 vs. Experiment 3). In contrast, the mean RTs to the English control items across the three experiments were affected much less by these changes. Similarly, De Groot et al. (2000) obtained very similar RTs for the English control conditions in English lexical decision without (Experiment 2) and with Dutch words (Experiment 3), and the same held for the Dutch control conditions in Dutch lexical decision without and with English words. Rather than explaining the small differences between the control conditions across the experiments of the two studies in terms of relative language activation, it is more elegant to explain them as small changes in the decision process applied in the different experiments or small variations in the participants' L2 proficiency. In sum, it does not seem to be possible to suppress the activation of one reading of an interlingual homograph even when that would undoubtedly improve performance.

However, one could argue that the alternative reading of interlingual homographs is always completely supported by input information. Perhaps the bottom-up support for these items is so strong that it overrides any attempt at top-down inhibition. This is not the case for neighbors of the target item, so conceivably for neighborhood materials non-target candidates could be more easily suppressed. However, cross-experimental comparisons with respect to neighborhood density suggest no suppression either. Van Heuven et al. (1998) had Dutch-English participants of one population perform a blocked or a mixed Progressive Demasking task. According to a relative language activation view, adding words from the non-target language should induce extra activation of the non-target language, which should affect especially the recognition of items in the weaker language (English for these bilinguals). However, RTs in the two experiments correlated .99 and .98 for English and Dutch target items (Dijkstra and Van Heuven, 1998, Tables 6.1
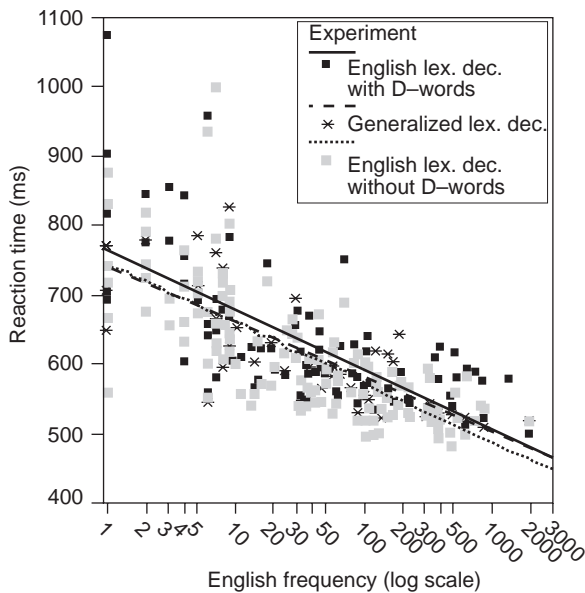
Figure 3. Regression lines reflecting the dependence of lexical decision RTs to English words of a considerable English frequency range on their frequency. These data were collected for Dutch-English students of psychology in visual lexical decision tasks including and excluding purely Dutch words, and a generalized Dutch-English lexical decision task (Schulpen et al., in preparation). There were no statistically significant differences between the lines in terms of their slopes.

and 6.2), so the result patterns were very much alike. The relative activation of the two languages was apparently not affected by mode of list presentation, because otherwise one would expect one of the two correlations to be much smaller (furthermore, regression analyses of the relationship between RTs and English or Dutch item frequency in the two experiments yielded nearly parallel lines).

Other evidence that items of the two languages are always fully activated is found in a recent study by Schulpen, Dijkstra and Schriefers (in preparation). This study replicated and extended the three experiments from Dijkstra et al. (1998(b)) with Dutch-English high school students of 15 and 17 years of age, students of psychology and Ph.D. students. Each experiment not only included interlingual homographs and matched English controls of different frequency groups, but also a set of exclusively English words that varied across a considerable frequency range. As before, Experiment 1 consisted of English words only, Experiment 2 also incorporated purely Dutch words that had to be rejected, and Experiment 3 contained English and Dutch words that all had to be accepted. Figure 3 depicts regression analyses indicating how RTs to English items of a considerable frequency range became

faster in all three experiments when English frequency increased. Just like in the earlier studies (Dijkstra et al., 1998(b); De Groot et al., 2000), the RTs for English target words presented in a list that also contained Dutch words (Experiment 2) were somewhat slower than when presented without Dutch words (Experiment 1), but the slope of the two regression lines was statistically identical. Schulpen et al. argue that this is evidence that the response to English words is affected only by a different decision process in Experiments 1 and 2. If task demands were to affect the relative language activation in the identification system, one would expect an interaction between English frequency and Experiment (1 or 2) because suppression should differentially affect the activation of lower and higher frequency items across experiments.

To conclude, the available evidence indicates that differences between these experiments, involving participants from the same population and the same stimulus materials, result from changes in decision parameter settings related to task demands rather than from changes in the relative activation levels of items of the two languages. This conclusion goes against such models as the original BIA model, the IC model and the theoretical view of the language mode. It is more compatible with the positions expressed by De Groot et al. (2000) and Dijkstra et al. (2000(a)), and will be taken into account during the future implementation of the task/decision level in the BIA+ model.

4.6 Relationship between word identification and task demands

The BIA+ distinction between an identification system and a task/decision system is in line with the view that in performing a task (such as lexical decision) an early preconscious, automatic level of processing may be followed by an attention-sensitive level in which percepts are selected with reference to contextual factors of various sorts and linked to particular responses relevant to the task at hand (cf. Altenberg and Cairns, 1983, p. 187; Dupoux and Mehler, 1992; Balota, Paul and Spieler, 1999).

The task schema specifies the series of mental processing steps to be taken or operations (the "algorithm") to be carried out to perform the specific task at hand (Green, 1986, 1998; Norman and Shallice, 1986). The task schema is set up during the practice set or retrieved from memory to perform the task. The decision mechanism is part of the task schema. It continuously reads out the activation in the identification system. At the appropriate cognitive step in the schema, it weighs different kinds of activation

input from the identification system relative to others to arrive at a response.

The decision will often depend on lexical selection, based on the crossing of an activation threshold for a lexical candidate. For instance, in lexical decision the input letter string activates orthographic, semantic and phonological codes, all of which could allow a discrimination of word and nonword input. However, code use might be quite different in language decision, in which participants press one button if a presented item belongs to one language (e.g., English) and another button if it belongs to another language (e.g., Dutch). In language decision, only codes that facilitate the retrieval of language membership information (language tags) are useful to arrive at a correct response.

Thus, different schemas underly different tasks, and sometimes one task may be performed by different schemas. Apart from the general organization of the task schema, performance may further be flexible in terms of the representations on which read-out is based and with respect to the amount of evidence required in favor of particular lexical options at different moments in time. For instance, as Grainger and Jacobs (1996) have argued, the composition of the stimulus list (e.g., the types of nonwords) may affect these decision criteria (also see Lima and Huntsman, 1997; Lupker, Brown and Colombo, 1997; Taylor and Lupker, 2001).

After the task schema is set up, dynamic adaptations during the experiment usually concern a fine-tuning of parameters. On the basis of the practice set, recognition thresholds and temporal deadlines are set that may be adapted during the experiment dependent on local stimulus list composition. For instance, the threshold on item identification is context-sensitive and may be adapted after each trial; the temporal deadline may be set depending on the amount of activation early in processing or the frequency of the previous items in the list, etc.

Context-sensitive adaptations of decision parameter settings can take place during stimulus list presentation without conscious intervention of the participant. For instance, the selection threshold for words may be lowered if the nonwords in the list are easy to reject and response speed can be increased at an acceptable cost in error rate. As another example, in Dutch lexical decision the threshold for saying "yes" to Dutch words may be lowered when it turns out that all English words present are low frequency and only few errors are made. The decision mechanism tries to optimize performance in terms of an internally set optimality criterion involving speed and accuracy (derived from the written, verbal and non-verbal instructions received from the experimenter).

A structural refiguration of the task schema may occur if the participant becomes aware of particular correlations between variables in the experiment. For example, Von Studnitz and Green (2002) have proposed that a "conjoint test" may develop in language-specific lexical decision if participants realize that both readings of interlingual homographs could be used for responding when there are no context words belonging exclusively to the non-target language. The conjoint test uses the simultaneous presence of lexical activation in two languages to speed up the response relative to one-language control items.

We further propose that the identification and task/decision systems, though interconnected, are partially independent. The identification system provides continuous output to the task/decision system, which is likely to use that output for response selection. Nevertheless, the task/decision system to some extent functions on its own. First, a word presented to the identification system may be recognized some time before a decision can be made, because certain information necessary to perform the task (e.g., language information in language decision) is not yet available or because the participant's strategy is to wait for additional information. Second, the opposite situation can also arise, for instance, in lexical decision. A participant may respond before lexical identification has taken place or if only limited information is available (e.g., orthographic but no semantic or phonological information). Some approaches have called this type of strategy "sophisticated guessing" because a decision may be made on essentially too little information to ensure a correct response in all cases. In the Multiple Read Out model (MROM) for lexical decision by Grainger and Jacobs (1996), the sum or sigma criterion for responding is part of a similar mechanism. Thus, the decision mechanism uses the activation within the identification system as a basis for decision, but it need not await the settling of the activation process, and it could, if it is set that way, also use other sources of information (e.g., guessing on the basis of stimulus list composition or participant expectancies).

We assume that the two systems have their own criteria for action triggering (i.e., lexical selection and response selection/execution). The identification system can be said to recognize a word (i.e., selects a single lexical candidate with enough certainty) when the system reaches a stable state. The task/decision system triggers a response if its own criteria are fulfilled, some of which pertain to lexical activation, while others refer to optimization criteria. The decision mechanism estimates the chances of making an error if a response were initiated at a particular

moment in time. This does not necessarily concern only the chance of giving a wrong response (e.g., in lexical decision), but can also pertain to the moment in time one starts speaking during word naming. For instance, starting to speak before the complete articulatory representation is available might interrupt the naming response. The task/decision mechanism must dynamically evaluate the perceived activation in different parts of the identification system and bind it to a particular response in such a way that optimal performance is likely to arise.

Because participants can only control the operation of the task/decision mechanism, they have only a specific type of control over their performance. After a task schema is set up, particular to the experimental situation, performance can only be adapted by resorting to strategies involving selective read-out, dynamic adjustment of identification criteria, and so on, but not by adapting the activation level of individual items or languages. For instance, Van Heuven et al. (1998) observed that Dutch/English high-proficiency participants seemed to modulate the relative contribution of English neighbors to Dutch word identification times in the course of the blocked progressive demasking experiment. Cross-linguistic effects changed from clear inhibition early in the experiment to much smaller effects later. According to the present model, both English and Dutch lexical candidates were always activated, but participants adapted their decision criteria to optimize their performance.

4.7 Stimulus-response binding in lexical decision

The task/decision system specifies how activated and selected representations in the identification system are bound to possible responses. For instance, in word naming the selected phonological and articulatory representations are used as a basis for response. In lexical decision, a "yes" response is triggered by the presence of sufficient lexical activation at a particular moment in time, while a "no" response is given if lexical activation is low when a temporal deadline is reached or when a mismatch is detected between the most activated lexical candidate and the input string at some critical evaluation moment. Here the schema might opt to use information from different sources in parallel, but presently available evidence suggests that orthographic representations play a major role (Pexman and Lupker, 2001). In generalized lexical decision, the decision mechanism is slightly more complex because the production rule is to "respond 'yes' if the input is a word in either L1 or L2". For language decision, responses ("press the English or Dutch response button") are assumed

to be bound to language tag representations, connected to language specific form representations or lemmas.

In some experimental paradigms, stimulus-response binding might already be partially prepared at an early stage. That could happen, for instance, in language decision, where the language tag for all Dutch words could be bound to the "left" response and the tag for English words to the "right" response even before the experiment. This would explain why the RTs in that task are not so much slower than in lexical decision. (An interesting possibility is that "global activation" effects arise because nonwords or words have many neighbors that are connected to the same response.)

The view of a dynamic adaptation of stimulus-response bindings and decision criteria may help to explain some remarkable lexical decision results in the literature, for instance, the null-results observed for interlingual homographs in De Moor (1998), Dijkstra et al. (1998(b), 2000(a)) and De Groot et al. (2000), or the results obtained by Dijkstra et al. (2000(a)) in the first and second part of their study (see section 4.5 above). Suppose that bilingual participants in these English lexical decision experiments try to follow their task instruction to the letter and set up a task schema in which they bind the "yes" response strongly to all English words, using the language tag for this purpose. They also understand in a more abstract way that other words and nonwords must receive a "no" response. Now suppose that an interlingual homograph is presented. As long as there is no strong binding of the Dutch reading of the homograph to the "no" response, there will be no strong response competition (even though both readings of the homograph are activated). This will lead to the absence of RT differences between homographs and one-language control words. As soon as a Dutch word is presented in the stimulus list (or practice list), however, the response binding of Dutch stimuli to the "no" response will be strengthened considerably, leading to much stronger response competition effects and inhibition for the interlingual homographs. This is in correspondence with the observed results. Furthermore, in the study by Dijkstra et al. (2000(a)), English control words were not affected by the transition from the first part of the experiment to the second because they were bound to "yes" only. Nonwords were not affected by this change either, because they were bound to the "no" response in a completely different way, namely in terms of the absence of enough lexical evidence at a temporal deadline.

So far this account is based upon decision level effects and it does not account for cross-linguistic

effects found in comparable studies. As noted in section 4.1, Dijkstra et al. (1999) did observe significant RT-differences to interlingual homographs (and cognates) and controls for Dutch-English bilinguals who performed an English lexical decision task. In the BIA+ model, such effects originate from interactions in the mental lexicon before the response comes about. Note that until the target word is identified and its language tag retrieved for responding, cross-linguistic interactions can arise within the mental lexicon. In language specific lexical decision (e.g., English), these effects may not be very visible, because the non-target language (e.g., Dutch) is not directly relevant to the response. However, when the stimulus list is mixed and the other-language words are connected to a "no" response, this response binding serves as a magnifying glass and strong inhibition effects become visible.

In this view, there is an important difference in the nature of the observed relative frequency effects and the cross-linguistic code similarity effects (Dijkstra et al., 1998(b), 1999). Word frequency is a determinant of the speed with which lexical information can be retrieved, whereas the size of code similarity effects depends, for instance, on the time interval between stimulus presentation and response selection during which interactions within the lexicon can take place.

We conclude that the reviewed lexical decision studies were affected by intra-lexical effects (explaining, for instance, effects of cross-linguistic similarity and relative frequency) as well as extra-lexical effects at the task/decision level (explaining, for instance, the change of effects from part 1 to part 2 in the Dijkstra et al., 2000, study). Implementing stimulus-response binding effects in the BIA+ model will constitute a challenging first step in the direction of implementing task schemas.

4.8 Stimulus-response binding in language switching

Several studies have investigated effects of task and language switching (Grainger and Beauvillain, 1987; Von Studnitz and Green, 1997; Meuter and Allport, 1999; Thomas and Allport, 2000), providing important information about the relationship between identification and task /decision system. Green (1998) proposed two possible loci of language switch costs in bilingual processing, namely within the lexico-semantic system and at the task schema level (Green, 1998, p. 73). In the BIA model (Dijkstra and Van Heuven, 1998) these two loci were confounded because the language nodes had both representational and functional significance. Research by Von Studnitz and Green (1997) and by Thomas and

Allport (2000) indicates that the most important source of switch costs in most processing tasks probably does not lie within the lexicon.

The BIA+ model attempts to account for cross-trial findings in stimulus lists by distinguishing three components to language switching effects: the task and its associated response bindings (e.g., "if word, press 'yes' button"), the actual response (e.g., "yes" or "no"; naming response), and item language (e.g., English or Dutch). Thus, BIA+ does not assume that lexical decision requires the activation and competition of more than one task schema, as proposed by Thomas and Allport, 2000). In generalized lexical decision, the actual response ("yes") does not change when the language of the next word changes relative to a previous word. This might suggest a lexical locus of the language switching effect in this task, but decision level adaptations are also possible, for instance, (1) between-trial adaptations of the relative recognition threshold for English and Dutch words; (2) in terms of the attention switch in verifying the connection between, e.g., an English word to the "yes" response to verifying that between the following Dutch word and the "yes" response. Furthermore, in the standard language specific lexical decision task involving words from two languages (e.g., Dijkstra et al., 1998(b)), a language switch implies a change in language and actual response (e.g., from "yes" to "no"), but not in response binding. In contrast, in the alternate runs variant of language specific lexical decision, a switch between trials implies a change in all three factors. Possibly, the decision system is sensitive to the presence of matches and mismatches between the three types of factors (cf. the frequent errors in the go/no-go task by Dijkstra et al., 2000(b), where the detection of a word rather than its language information led to a response, see section 2.5 above).

5 Generalizability of the BIA+ model

The BIA+ model is a model for bilingual visual word recognition. However, we think that it can be generalized to other domains in a number of respects. For instance, we assume that the distinction between word retrieval and task/decision system will also be valid for bilingual auditory word recognition and bilingual word production (cf. Green, 1998). Furthermore, the available evidence from these two other domains suggests that there may be similarities in terms of lexical access mechanisms as well. For lack of space, we refer the reader to Kroll and Dijkstra (2002) for a discussion of the similarities between bilingual language production and bilingual visual word recognition in terms of the BIA model. Here we

will compare bilingual visual and auditory word recognition in some detail.

Just like in the visual domain, the presentation of a spoken word is assumed to lead to the activation of a number of lexical candidates (for overviews see Frauenfelder, 1996; Dijkstra, in press). One of the available models of spoken word recognition, the TRACE model (McClelland and Elman, 1986), incorporates a theoretical view of auditory word recognition that is directly analogous to that of its visual counterpart, the Interactive Activation model (also see Grosjean, 1997, for BIMOLA, a bilingual model for auditory word recognition). However, in auditory models the nature of the set of lexical competitors is generally different from that in the visual modality. In spoken word recognition, it is the initial part of a word that leads to the activation of a set of matching lexical candidates. The activated collection of possible competing word candidates is often called the (word-initial) *cohort*. This collection of possible words is assumed to depend on the phonemes it consists of (just like the letters in visual word recognition), or perhaps on a matrix of activated phonetic/phonological features. Later mismatching information may lead to the exclusion of word candidates (and according to some models later matching information may activate possible candidates at a delay). Somewhere in the signal (possibly before the end of the input has been reached), enough information is available to uniquely identify the presented stimulus word and reject other possible word candidates with (enough) certainty.

Given the similarity between visual and auditory word recognition in terms of the parallel activation of lexical competitors, it is reasonable to assume that there is language non-selective access in the auditory modality as well. However, the auditory signal contains language-specific sublexical cues that have no direct counterpart in the visual modality when words are read from two languages with the same script. For instance, there are differences in phonemic repertoires across languages that can potentially be used to reduce the cohort set already early during bilingual auditory word recognition.

Surprisingly enough, most of the few studies that have been done indicate that auditory lexical access is language non-selective in nature, even though participants are at the same time sensitive to sublexical cues (Li, 1996; Marian and Spivey, 1999; Schulpen, Dijkstra, Schriefers and Hasper, submitted; but see Pallier, Colomé and Sebastián-Gallès, 2001, for a different view). For instance, in three experiments, Schulpen et al. (submitted) investigated how interlingual homophones are recognized during bilingual spoken word recognition. In a gating study, word

identification and language membership decisions by Dutch-English bilinguals were found to be slower for interlingual homophones than for one-language control items. At the same time, participants were sensitive to subphonemic cues when they assessed the language a word belonged to. These findings suggest that auditory lexical access is language non-selective while sensitive to language specific characteristics of the input. Two additional on-line experiments tested this view using a cross-modal priming paradigm. Visual targets included interlingual homophones and one-language control items. Auditory distractors differed in their amount of form overlap with the visual targets. Visual lexical decisions were fastest for one-language controls preceded by their auditory equivalents (e.g., the English auditory prime /frame/ followed by the related visual target word FRAME). Responses to interlingual homophones accompanied by their English or Dutch distractor version (e.g., the visually presented English target word LEAF preceded by the auditory English prime /leaf/ or the Dutch prime /lief/, meaning "dear") were also faster than all unrelated conditions. However, they were slower than the related one-language control conditions, providing evidence for on-line competition of their two readings. Different effects of English and Dutch distractors on homophone processing in Experiment 3 suggest that participants used sublexical cues to differentiate the two readings of a homophone after language non-selective access.

Apparently, a slight mismatch is not sufficient to exclude word candidates from another language from the cohort (Marian and Spivey, 1999; Schulpen et al., submitted). Future studies will have to determine how much mismatching information is necessary to induce language-selective access results and how it affects the time-course of activation of L1 and L2 candidates. One would expect that sublexical and other cross-linguistic cues (e.g., in terms of word stress and tonal information) must somehow be able to help to reduce the number of possible competitors from another language under some circumstances.

In sum, in spite of considerable differences in terms of representations and processes, bilingual visual and auditory word recognition both appear to entail a language non-selective access procedure involving an initial activation of word candidates from different languages. Even though modality specific processing characteristics are to be expected, so far most evidence seems to be compatible with the BIA+ framework.

6 Conclusion

In this paper we have presented a new theoretical

framework for bilingual word recognition, called the BIA+ model, which consists of an identification system that provides information about activated representations from different languages to a decision and response selection mechanism operating as part of a task schema. Although related to the BIA model and the IC model, the BIA+ model is more explicit than other models with respect to the time-course of the bilingual word identification process, the interactions between different types of representations (orthographic, phonological, semantic) and the representation of interlingual homographs and language membership tags. Furthermore, the model makes a distinction between the effects of linguistic and non-linguistic context on performance. Linguistic context effects (e.g., sentence context) can directly affect the activity in the word identification system, while non-linguistic context effects (e.g., participant strategies) can only affect the task/decision system. These distinctions lead to many testable model predictions, just like the model's detailed account on stimulus-response binding in lexical decision and in language switching experiments.

Evidence supporting the distinction between the word identification and the task/decision system was derived in part from a comparison of closely related variants of the lexical decision task. A comparison of the result patterns of language-specific and generalized task variants including stimuli from one or two languages was helpful in demonstrating that task demands can systematically modify the output pattern of the underlying word identification system. Evidence from other techniques, such as progressive demasking, language decision and word naming (and from ''classical'' paradigms such as recognition and recall in pure and mixed lists) is compatible with these conclusions. Thus, although different types of tasks necessarily require different ways of arriving at a response, they all appear to reflect a common underlying pattern of activation in the word identification system.

At present, the BIA+ model is only implemented with respect to the orthographic representations in the identification system (allowing the simulations in sections 2.3, 2.4, and 2.6). This is the part the model shares with the BIA model. An important difference between the BIA model and the BIA+ model is that the latter model rejects the powerful asymmetic top-down inhibition mechanism from language nodes to word nodes that has accounted for shifts in result patterns within and between experiments. The task/decision component is proposed as a replacement mechanism that should perform the same functions. Additional empirical data must be collected to test the assumption of BIA+ that the task/decision system cannot affect relative language activation in the identification system. Filling many more of the remaining gaps in our knowledge about bilingual lexical processing is necessary before a complete implementation of the BIA+ or of any other bilingual word recognition model can be built.

References

Altarriba, J., Kroll, J. F., Sholl, A. & Rayner, K. (1996). The influence of lexical and conceptual constraints on reading mixed-language sentences: Evidence from eye fixations and naming times. *Memory & Cognition*, 24, 477–492.

Altenberg, E. P. & Cairns, H. S. (1983). The effects of phonotactic constraints on lexical processing in bilingual and monolingual subjects. *Journal of Verbal Learning and Verbal Behavior*, 22, 174–188.

Andrews, S. (1989). Frequency and neighborhood effects on lexical access: Activation or search? *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 802–814.

Balota, D. A., Paul, S. T. & Spieler, D. H. (1999). Attentional control of lexical processing pathways during word recognition and reading. In S. Garrod & M. Pickering (eds.), *Language processing*, pp. 15–57. Hove: Psychology Press.

Bijeljac-Babic, R., Biardeau, A. & Grainger, J. (1997). Masked orthographic priming in bilingual word recognition. *Memory & Cognition*, 25, 447–457.

Coltheart, M., Davelaar, E., Jonasson, J. T. & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (ed.), *Attention and performance VI*, pp. 535–555. New York: Academic Press.

Cristoffanini, P., Kirshner, K. & Milech, D. (1986). Bilingual lexical representation: The status of Spanish-English cognates. *Quarterly Journal of Experimental Psychology*, 38A, 367–393.

De Groot, A. M. B., Delmaar, P. & Lupker, S. J. (2000). The processing of interlexical homographs in a bilingual and a monolingual task: Support for non-selective access to bilingual memory. *Quarterly Journal of Experimental Psychology*, 53, 397–428.

De Moor, W. (1998). Visuele woordherkenning bij tweetalige personen [Visual word recognition in bilinguals]. Master thesis, University of Ghent.

Dijkstra, A. (Ton) (1998). From tag to task: Coming to grips with bilingual control issues. *Bilingualism: Language and Cognition*, 1, 51–66.

Dijkstra, A. (Ton) (in press). Word recognition and lexical access II: Connectionist approaches. In D. A. Cruse, F. Hundsnurscher, M. Job & P. R. Lutzeier (eds.), *Lexikologie–Lexicology*. Berlin: Walter de Gruyter. (Article 265.)

Dijkstra, A. (Ton), De Bruijn, E., Schriefers, H. J. & Ten Brinke, S. (2000(a)). More on interlingual homograph recognition: Language intermixing versus explicitness of instruction. *Bilingualism: Language and Cognition*, 3, 69–78.

Dijkstra, A. (Ton), Grainger, J. & Van Heuven, W. J. B. (1999). Recognition of cognates and interlingual homographs: The neglected role of phonology. *Journal of Memory and Language*, 41, 496–518.

Dijkstra, A. (Ton) & Kempen, G. (1984). *Taal in uitvoering* [Language in progress]. Groningen: Wolters-Noordhoff.

Dijkstra, A. (Ton), Timmermans, M. & Schriefers, H. (2000(b)). Cross-language effects on bilingual homograph recognition. *Journal of Memory and Language*, 42, 445–464.

Dijkstra, A. (Ton) & Van Heuven, W. J. B. (1998). The BIA model and bilingual word recognition. In J. Grainger & A. M. Jacobs (eds.), *Localist connectionist approaches to human cognition*, pp. 189–225. Mahwah, NJ: Lawrence Erlbaum Associates.

Dijkstra, A. (Ton), Van Heuven, W. J. B. & Grainger, J. (1998(a)). Simulating competitor effects with the Bilingual Interactive Activation Model. *Psychologica Belgica*, 38, 177–196.

Dijkstra, A. (Ton), Van Jaarsveld, H. & Ten Brinke, S. (1998(b)). Interlingual homograph recognition: Effects of task demands and language intermixing. *Bilingualism: Language and Cognition*, 1, 51–66.

Dupoux, E. & Mehler, J. (1992). Unifying awareness and on-line studies of speech: A tentative framework. In J. Alegria, D. Holender, J. Junca de Morais & M. Radeau (eds.), *Analytic approaches to human cognition*, pp. 59–75. Amsterdam: Elsevier Science Publishers.

Duffy, S., Morris, R. & Rayner, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory and Language*, 27, 429–446.

Elston-Güttler, K. & Williams, J. N. (submitted). Bilingual ambiguities: Sentence context effects on processing interlingual homographs and translated L1 homographs in the L2.

Ferrand, L. & Grainger, J. (1993). The time-course of orthographic and phonological code activation in the early phases of visual word recognition. *Bulletin of the Psychonomic Society*, 31, 119–122.

Font, N. (2001). Rôle de la langue dans l'accès au lexique chez les bilingues: Influence de la proximité orthographique et sémantique interlangue sur la reconnaissance visuelle de mots. Ph.D. thesis, Université Paul Valery, Montpellier.

Frauenfelder, U. (1996). Computational models of spoken word recognition. In A. Dijkstra & K. De Smedt (eds.), *Computational psycholinguistics: AI and connectionist models of human language processing*, pp. 114–138. London: Taylor & Francis.

Gerard, L. D. & Scarborough, D. L. (1989). Language-specific lexical access of homographs by bilinguals. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 305–313.

Grainger, J. & Beauvillain, C. (1987). Language blocking and lexical access in bilinguals. *Quarterly Journal of Experimental Psychology*, 39A, 295–319.

Grainger, J. & Dijkstra, A. (Ton) (1992). On the representation and use of language information in bilinguals.

In R. J. Harris (ed.), *Cognitive processing in bilinguals*, pp. 207–220. Amsterdam: Elsevier.

Grainger, J. & Dijkstra, A. (Ton) (1996). Visual word recognition. In A. (Ton) Dijkstra & K. De Smedt (eds.), *Computational psycholinguistics: AI and connectionist models of human language processing*, pp. 139–165. London: Taylor & Francis.

Grainger, J. & Jacobs, A. M. (1996). Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review*, 103, 518–565.

Grainger, J. & O'Regan, J. K. (1992). A psychophysical investigation of language priming effects in two English-French bilinguals. *European Journal of Cognitive Psychology*, 4, 323–339.

Grainger, J., O'Regan, J. K., Jacobs, A.M. & Segui, J. (1989). On the role of competing word units in visual word recognition: The neighborhood frequency effect. *Perception & Psychophysics*, 45, 189–195.

Grainger, J. & Segui, J. (1990). Neighborhood frequency effects in visual word recognition: A comparison of lexical decision and masked identification latencies. *Perception & Psychophysics*, 47, 191–198.

Green, D. W. (1986). Control, activation, and resource: A framework and a model for the control of speech in bilinguals. *Brain & Language*, 27, 210–223.

Green, D. W. (1998). Mental control of the bilingual lexico-semantic system. *Bilingualism: Language and Cognition*, 1, 67–81.

Grosjean, F. (1997). Processing mixed language: Issues, findings and models. In A. M. B. De Groot & J. F. Kroll (eds.), *Tutorials in bilingualism: Psycholinguistic perspectives*, pp. 225–254. Hillsdale, NJ: Lawrence Erlbaum Associates.

Jared, D. & Kroll, J. F. (2001). Do bilinguals activate phonological representations in one or both of their languages when naming words? *Journal of Memory and Language*, 44, 2–31.

Kroll, J. F. & Dijkstra, A. (Ton) (2002). The bilingual lexicon. In R. Kaplan (ed.), *Handbook of applied linguistics*, pp. 301–321. Oxford: Oxford University Press.

Lalor, E. & Kirsner, K. (2000). Cross-lingual transfer effects between English and Italian cognates and noncognates. *International Journal of Bilingualism*, 4, 385–398.

Lemhöfer, K. & Dijkstra, A. (Ton) (submitted). Recognizing cognates and interlingual homographs: Effects of code similarity in generalized lexical decision.

Levelt, W. J. M., Roelofs, A. & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, 21, 1–38.

Li, P. (1996). Spoken word recognition of code-switched words by Chinese-English bilinguals. *Journal of Memory and Language*, 35, 757–774.

Lima, S. D. & Huntsman, L. A. (1997). Sequential dependencies in the lexical decision task. *Psychological Research*, 60, 264–269.

Lucas, M. (1999). Context effects in lexical access: A meta-analysis. *Memory and Cognition*, 27, 385–398.

Lupker, S. J., Brown, P. & Colombo, L. (1997). Strategic

control in a naming task: Changing routes or changing deadlines? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 570–590.

Marian, V. & Spivey, M. (1999). Activation of Russian and English cohorts during bilingual spoken word recognition. In M. Hahn & S. C. Stoness (eds.), *Proceedings of the Twenty-first Annual Conference of the Cognitive Science Society*, pp. 349–354. Mahwah, NJ: Lawrence Erlbaum.

Mathey, S. & Zagar, D. (2000). The neighborhood distribution effect in visual word recognition: Words with single and twin neighbors. *Journal of Experimental Psychology: Human Perception and Performance*, 26, 184–205.

McClelland, J. L. & Elman, J. (1986). Interactive processes in speech perception: The TRACE model. In J. L. McClelland, D. E. Rumelhart and the PDP Research Group (eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (vol. II), pp. 58–121. Cambridge, MA: MIT Press.

McClelland, J. L. & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception, Part 1: An account of basic findings. *Psychological Review*, 88, 375–405.

McCormack, P. D. (1977). Bilingual linguistic memory: The independence-interdependence issue revisited. In P. A. Hornby (ed.), *Bilingualism: Psychological, social, and educational implications*, pp. 57–65. New York: Academic Press.

Meuter, R. F. I. & Allport, A. (1999). Bilingual language switching in naming: Asymmetrical costs of language selection. *Journal of Memory and Language*, 40, 25–40.

Norman, D. A. & Shallice, T. (1986). Attention to action: Willed and automatic control of behaviour. In R. J. Davidson, G. E. Schwartz & D. Shapiro (eds.), *Consciousness & self-regulation* (vol. 4), pp. 1–18. New York: Plenum Press.

Pallier, Ch., Colomé, A. & Sebastián-Gallès, N. (2001). The influence of native-language phonology on lexical access: Exemplar-based versus abstract lexical entries. *Psychological Science*, 12, 445–449.

Pecher, D. (2001). Perception is a two-way junction: Feedback semantics in word recognition. *Psychonomic Bulletin & Review*, 8, 545–551.

Pexman, P. M. & Lupker, S. J. (2001). Ambiguity and visual word recognition: Can feedback explain both homophone and polysemy effects? *Canadian Journal of Experimental Psychology*, 53, 323–334.

Plaut, D. C., McClelland, J., Seidenberg, M. S. & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56–115.

Raab, D. H. (1962). Statistical facilitation of simple reaction time. *Transactions of the New York Academy of Sciences*, 24, 574–590.

Reimer, J. F., Brown, J. S. & Lorsbach, T. C. (2001). Orthographically mediated inhibition effects: Evidence of activational feedback during visual word recognition. *Psychonomic Bulletin & Review*, 8, 102–110.

Sánchez-Casas, R., Davis, C. W. & García-Albea, J. E. (1992). Bilingual lexical processing: Exploring the cognate/non-cognate distinction. *European Journal of Cognitive Psychology*, 4, 311–322.

Schulpen, B., Dijkstra, A. (Ton) & Schriefers, H. J. (in preparation). A cross-sectional study on the development of automatization and cognitive control in Dutch-English bilinguals, part 1. [working title]

Schulpen, B., Dijkstra, A. (Ton), Schriefers, H. J. & Hasper, M. (submitted). Recognition of interlingual homophones in bilingual auditory word recognition.

Snodgrass, J. G. & Mintzer, M. (1993). Neighborhood effects in visual word recognition: Facilitatory or inhibitory? *Memory & Cognition*, 21, 247–266.

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

Tabossi, P. & Zardon, F. (1993). Processing ambiguous words in context. *Journal of Memory and Language*, 32, 359–372.

Tanenhaus, M. K., Leiman, J. M. & Seidenberg, M. S. (1979). Evidence for multiple stages in the processing of ambiguous words in syntactic context. *Journal of Verbal Learning and Verbal Behaviour*, 18, 427–441.

Taylor, T. E. & Lupker, S. J. (2001). Sequential effects in naming: A time-criterion account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 117–138.

Thomas, M. S. C. & Allport, A. (2000). Language switching costs in bilingual visual word recognition. *Journal of Memory and Language*, 43, 44–66.

Van Hell, J. G. & Dijkstra, A. (Ton) (in press). Foreign language knowledge can influence native language performance in exclusively native contexts. *Psychonomic Bulletin & Review*.

Van Heste, T. (1999). Visuele woordherkenning bij tweetaligen [Visual word recognition in bilinguals]. Master thesis, University of Leuven.

Van Heuven, W. J. B. (2000). Visual word recognition in monolingual and bilingual readers: Experiments and computational modeling. Ph.D. Thesis, University of Nijmegen. (NICI Technical Report 20–01.)

Van Heuven, W. J. B., Dijkstra, A. (Ton) & Grainger, J. (1998). Orthographic neighborhood effects in bilingual word recognition. *Journal of Memory and Language*, 39, 458–483.

Van Heuven, W. J. B., Dijkstra, A. (Ton), Grainger, J. & Schriefers, H. (2001). Shared neighborhood effects in masked orthographic priming. *Psychonomic Bulletin & Review*, 8, 96–101.

Von Studnitz, R. E. & Green, D. W. (1997). Lexical decision and language switching. *International Journal of Bilingualism*, 1, 3–24.

Von Studnitz, R. E. & Green, D. W. (2002). Interlingual homograph interference in German-English bilinguals: Its modulation and locus of control. *Bilingualism: Language and Cognition*, 5, 1–23.

PEER COMMENTARIES

On the temporal delay assumption and the impact of non-linguistic context effects

MARC BRYSBAERT
Royal Holloway, University of London,
Department of Psychology, Egham TW20 0EX, UK.
E-mail: *marc.brysbaert@rhul.ac.uk*
ILSE VAN WIJNENDAELE
Universiteit Leuven
WOUTER DUYCK
Universiteit Gent

It is not easy to comment on Dijkstra and Van Heuven's model because there are many more aspects we agree with than aspects we feel uncomfortable about. Indeed, the BIA model has played an enormous role in showing us how bilingual visual word recognition can be achieved without recurrence to the intuitively appealing – but wrong – ideas of separate, language-specific lexicons and language-selective access. As in many other research areas, a working computational model has been much more influential in convincing critical readers (and researchers) than any series of empirical findings. The BIA+ model inherits this strength and, hopefully, in the coming years will be implemented in enough detail to exceed its predecessor. In the rest of this comment, we would like to put a cautionary note behind the temporal delay assumption introduced in the target article and provide some additional corroborating evidence for the lack of non-linguistic effects on early processes in the identification system.

The "temporal delay assumption" claims that L2 phonological and semantic codes are delayed in activation relative to L1 codes. This seems to imply that the delay will be substantial and constant for all types of words (only depending on the proficiency of the bilingual). With respect to the orthography-phonology conversion, we have quite strong evidence that this is not the case. If the activation of phonology in L2 were rather slow (and/or weak), then one would expect less impact of phonology in second language processing than in first language processing. Probably the best way to study the importance of phonological coding in visual word recognition is to make use of the masked priming paradigm. In this paradigm, a target word (e.g., "side") is preceded by a tachistoscopically presented homophonic non-word prime that sounds the same (e.g., "syed") or by a graphemic control prime that shares the same number of letters but does not sound the same (e.g., "soed"). Evidence of phonological encoding is obtained when target word processing is better after the homophonic prime than after the graphemic control prime. Such a homophonic priming effect has been reported in many languages including English, Dutch and French (see Brysbaert, 2001 for a review).

In a series of experiments using Dutch and French target words, we showed that the magnitude of the homophonic priming effect was equally large in L2 as in L1 (Van Dyck and Van de Poel, 1999; Brysbaert, Van Wijnendaele and Brysbaert, 2002; Brysbaert, in press), suggesting that the activation of phonological information from a visually presented word is equally strong in both languages. So, the

French word "faim" [hunger] had more chances of being recognised after the homophonic prime "fain" than after the graphemic control prime "faic", both for French-Dutch bilinguals and for Dutch-French bilinguals. The same effect was found with Dutch target words like "fijn" [fine] preceded by the primes "fein" and "foun". These studies made use of bilinguals with a clear dominance of the first language (i.e., non-balanced bilinguals).

Apart from the above empirical findings, there are also theoretical reasons why the activation of L2 phonology need not always trail behind the L1 activation. When two languages share the same alphabet, they are also likely to have many similar letter-sound correspondences, in particular for the consonants. For instance, the letters "b", "c" and "d" stand for very much the same sounds in Dutch and French, and there are no reasons to believe why L2 processing would not make use of these already existing, compatible L1 conversions. Furthermore, a number of letter signs (or letter combinations) may be unique to L2 and if these are encountered frequently enough, their cumulative frequency may outnumber the cumulative frequency of less common L1 conversions. Again, any learning-based system would predict that these L2 conversions are completed faster than the less frequent L1 conversions (unless both are at ceiling level, in which case conversion times will be the same). Finally, there are the interesting instances where the L2 letter-sound correspondences are incompatible with the existing L1 letter-sound correspondences. For instance, the grapheme "ee" is pronounced /i:/ in English (as in "meet") but /e:/ in Dutch (as in "leed" [sorrow]). Depending on the model of grapheme-phoneme conversions, one has to predict that such L2 conversions either do not happen (if the conversions are rule-based) or start to interfere with the original L1 conversions (as in connectionist-type learning models; see Van Wijnendaele and Brysbaert, 2002 for a detailed discussion). In the latter case, L2 proficiency will not only have implications for the processing of L2 but also for the processing of words in the mother tongue (Brysbaert, in press). Again, the end result is not simply a delay of L2 codes relative to L1 codes, but interactions between both types of codes that are much more dynamic than implied by the "temporal delay assumption".

Very much the same story could probably be told about the activation of semantics from lexical orthography and phonology. For those words that have exactly the same, well-defined meaning in L1 and L2, it is probably relatively easy to activate the existing patterns of semantic features from the new word forms. A typical class of such words are

the number words, which have exactly the same meaning in different languages. Tzelgov, Yehene, Kotler, and Alon (2000) showed that arbitrary symbols representing magnitudes after a short learning phase evoke the same semantic effects as the original Arabic digits. For instance, when digit pairs are shown in two fonts with different physical magnitudes, participants find it more difficult to indicate that 8 is the smaller symbol in the pair "2 – 8" than that 8 is the larger symbol in the pair "2 – 8". This effect is referred to as the Stroop-like size congruity effect. Tzelgov et al. observed that after a short training session the size congruity effect can be obtained with any pair of arbitrary symbols that have been introduced as "alternative digits". Duyck and Brysbaert (in preparation) wondered whether a similar effect would occur in number translation. First, they showed that bilinguals need less time to translate small verbal numerals (e.g. "two") than large numerals ("eight"), as could be expected on the basis of the finding that large numbers need more time to activate their semantic representation than small numbers (Brysbaert, 1995). Then, they trained a group of participants for half an hour on unknown, so-called "Estonian number words" and had them translate these words. Exactly the same number-magnitude effect was obtained as in the original study (which involved two well-known languages), strongly suggesting that the translation of the new Estonian words was semantically mediated as well. So, the activation of semantics from L2 words need not take longer than the activation from L1 words, provided that the meanings of both classes of words are the same and clearly specified. This will be particularly true for those words that in addition to full semantic overlap, also share form properties (i.e., cognates). Less rapid activation of semantics can be expected for L2 words that cannot profit from existing L1 word-meaning mappings (because they do not have literal translations) or that contradict the existing L1 mappings (as could be the case for homographs). Again, most learning algorithms would predict that the introduction of inconsistencies in the mappings due to the learning of a new language not only has implications for the speed of the L2 conversions, but is also likely to affect the existing L1 conversions.

The above paragraphs remind us that a verbal description of a computational model is not the same as a fully implemented and running model. Therefore, we are eager to see the future implementations of the BIA+ model. Our hunch is that the temporal delay assumption will be much more difficult to implement than the scenario we have outline above. In the remainder, we would like to finish on a more positive note, showing that the orthography-phonology conversion part of the word identification system is indeed impervious to the non-linguistic context.

In the priming studies described above, we not only looked at intra-language phonological priming but also at cross-language priming. In particular, we showed that not only is it possible to prime an L2 target word with an L2 homophonic prime, but also to prime an L2 target word with an L1 homophonic prime. So, for a Dutch-English bilingual (but not for an English monolingual), it is possible to prime the target word "blame" with the non-word "bleem", because "bleem" is a pseudohomophone of the target "blame" according to the Dutch spelling-sound correspondences. In addition, we showed that it is equally possible to prime an L1 target word with a L2 homophonic prime. So, the "bleem-blame" example also works with English-Dutch bilinguals. (Incidentally, this again shows that L2 phonology is not simply delayed relative to L1 phonology.)

Van Wijnendaele (2002, Chapter 7) asked whether the cross-lingual L2 on L1 phonological priming effect depended on the participants being aware of the bilingual nature of the study (as implied by Grosjean's "language mode"). To answer this question, she asked two groups of Dutch-French bilinguals (university students from Leuven) to take part in a Dutch word recognition experiment. The first group of students was told that the experiment was part of a bilingual word recognition study program and they completed a French word recognition experiment before embarking on "the Dutch part". The second group of students was simply told that the experimenter was interested in Dutch word recognition (i.e., the mother language of the students and also the language used at the university). The French language was not mentioned (as a matter of fact, special precautions were taken to make sure that the students could not find out that the experiment was part of a Ph.D. on bilingual word recognition).

The students were told that upper-case target words would be presented briefly on a computer screen between two rows of hash-signs (########). Their task was to try to recognise the word. Unknown to the students, the target words were preceded by a tachistoscopically presented, lower-case prime (presented too shortly to be visible). Half of the primes were Dutch non-words (homophonic or controls, such as "fein/foun" for the target word "FIJN" [nice]); half were French non-words (e.g. "deuque/delu" for the target word "DEUK" [dent]). The Dutch pseudo-homophones sounded like the target words if the Dutch letter-sound conversions were used; the French pseudo-homophones sounded like the target words if the French letter-sound conversions were used. The hypotheses were: (a) that the students would be more likely to identify the words after a homophonic prime than after a graphemic control prime, and (b) if Grosjean is right and if bilinguals are able to de-activate one part of their language system, we would not observe priming from French L2 primes on Dutch L1 words when participants are not aware of the fact that the French language matters for good performance (i.e. those participants in the monolingual mode). The results of the study are displayed in Table 1.

As can be seen in Table 1, there was no sign of any significant modulation of the phonological priming effect as a function of the language mode of the participants. As a matter of fact, the priming effect was slightly larger for the group of students who were not aware of the fact that knowledge of French mattered (the monolingual mode) than for the students who had been informed about the importance of the French language (the bilingual language mode). The only difference between the groups that could have some implications was the tendency of the participants in the bilingual mode to recognize fewer L1 target words

Table 1. *Percentage word identification as a function of language mode and prime type (for details, see Van Wijnendaele, 2002, Chapter 7).*

	Bilingual mode	Monolingual mode
deuque-DEUK	54%	65%
delu-DEUK	49%	58%
priming effect	5%	7%
fein-FIJN	71%	69%
foun-FIJN	62%	59%
priming effect	9%	10%

when they followed a L2 prime ("deuque/delu – DEUK") than when they followed an L1 prime ("fein/foun – FIJN"; see the first column of Table 1). No such difference was present for the participants in the monolingual mode. This seemed to indicate that there might be some inhibition from L2 on L1, if (a) students are in a bilingual language mode, and (b) there is bottom-up information in the prime strongly pointing to L2. However, two unpublished follow-up studies that tried to replicate this finding failed. Therefore, it is probably safe to assume, as Dijkstra and Van Heuven do, that the word identification system is impervious to non-linguistic context manipulations.

References

Brysbaert, M. (1995). Arabic number reading: On the nature of the numerical scale and the origin of phonological recoding. *Journal of Experimental Psychology: General*, 124, 434–452.

Brysbaert, M. (2001). Prelexical phonological coding of visual words in Dutch: Automatic after all. *Memory & Cognition*, 29, 765–773.

Brysbaert, M. (in press). Bilingual visual word recognition: Evidence from masked phonological priming. In S. Kinoshita & S. J. Lupker (eds.), *Masked priming: State-of-the-art*. Hove: Psychology Press.

Brysbaert, M., Van Dyck, G. & Van de Poel, M. (1999). Visual word recognition in bilinguals: Evidence from masked phonological priming. *Journal of Experimental Psychology-Human Perception and Performance*, 25, 137–148.

Duyck W. & Brysbaert, M. (in preparation). Forward and backward translation in balanced and unbalanced bilinguals requires conceptual mediation: the magnitude effect in number translation.

Tzelgov, J., Yehene, V., Kotler, L. & Alon, A. (2000). Automatic comparisons of artificial digits never compared: Learning linear ordering relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 103–120.

Van Wijnendaele, I. (2002). Cognitieve representatie van tweetaligheid: Visuele woordherkenning bij tweetaligen. Ph.D. thesis, Katholieke Universiteit Leuven, Department of Psychology.

Van Wijnendaele, I. & Brysbaert, M. (2002). Visual word recognition in bilinguals: Phonological priming from the second to the first language. *Journal of Experimental Psychology-Human Perception and Performance*, 3, 619–627.

The BIA++: Extending the BIA+ to a dynamical distributed connectionist framework

MAUD JACQUET
ROBERT M. FRENCH
Quantitative Psychology and Cognitive Science (B32), Department of Psychology, University of Liège, 4000 Liège, Belgium.
E-mail: *mjacquet@ulb.ac.be, rfrench@ulg.ac.be*

Dijkstra and van Heuven have made an admirable attempt to develop a new model of bilingual memory, the BIA+. Their article presents a clear and well-reasoned theoretical justification of their model, followed by a description of their model. The BIA+ is, as the name implies, an extension of the Bilingual Interactive Activation (BIA) model (Dijkstra and van Heuven, 1998; Van Heuven, Dijkstra and Grainger, 1998; etc), which was itself an adaptation to bilingual memory of McClelland and Rumelhart's (1981) Interactive Activation model of monolingual memory.

The authors provide a wealth of background on bilingual memory cross-lingual interference and priming effects in what amounts to a veritable review of the literature in this area. The model that they propose is designed to account for many of these empirically observed effects. In what follows we will center our discussion around three points related to the design of their model. These issues are:

- the use of modular vs. distributed representations;
- learning;
- emergence and self-organization of lexical items.

We will discuss each of these points in turn.

Overview

From the earliest days of the development of computer models of human cognitive capacities, one of the most significant problems that hung over the entire endeavor was the problem of hand-coded representations, or the Problem of Representation, as it is sometimes called (see, for example, Chalmers, French and Hofstadter, 1992; Elman, 1995). Over and over again, exaggerated claims were made for programs that supposedly discovered new mathematical theorems, solved complex problems, made scientific discoveries, took creative leaps or discovered analogies. And, each time, upon closer inspection, one discovered that the real reason these programs were able to do anything at all was because they had been given input data that had been carefully tailored so that the desired solution was, if not necessarily inevitable, at least not too difficult to produce. This does not, of course, mean that one can *never* use hand-coded representations. Obviously, at some point, modelers have to make decisions about the form of the input data that will inevitably influence the output of their programs. So the real issue is this: To what extent does the fit of a program's output to empirical data depend on the way in which the raw data was "pre-processed" (i.e., by a human) before being presented to the program. This will become a crucial issue in the discussion that follows.

Modular vs. distributed representations

We are by no means eliminativists (e.g., Churchland, 1995) who decry any use of any modular structures in modeling cognition. The need for modularity in programs simulating complex cognitive abilities is obvious. On the other hand, the addition of arbitrary processing modules every time new data conflicts with a previous "boxological" structure is not acceptable either. We believe that the necessary modules should be able to be explained *as an emergent by-product of the architecture*. This allows modelers to make use of modularity when necessary but only insofar as it can be explained as an emergent product of lower level mechanisms.

This leads us to our first point of contention with the BIA+. This model is a far cry from the old Information Processing models of the pre-connectionist era. In keeping with the philosophy of McClelland and Rumelhart's 1981 IA model, the letter level in the BIA+ emerges from the primitive-feature level and the word level from the letter level. But then the problems begin. The letter level gives rise to the word level but, unlike the monolingual IA case, at the word level *the lexical items are separated into their respective languages*. We are never told how this might work. Of course, if the program already knows that a word belongs to a particular language, then the job is much easier. But that's cheating. Little children learning two languages are rarely told explicitly that "chien" is a word in French, while "bug" is a word in English. They simply hear the two words and they gradually *learn* where they belong. So, the question for Dijkstra and van Heuven is: How could these two "language modules" (Dutch and English) have come about in the first place? What would "gate" a particular word to incorporate it in the Dutch module and how would an English word be incorporated into the English module? And, most seriously of all, what do the "language nodes" associated with each module correspond to? In a standard spreading-activation account with distributed internal representations, these language nodes do not strike us as necessary. People know they are writing in English (as opposed to, say, French) because a coherent, highly-interconnected set of representations of "English" items are currently active in their brains, period. There is no need to have an additional, explicit "language" node continually reminding the writer that he or she is writing in English.

Further, in our reading of their description of the BIA+, explicit language nodes would seem to pose problems that would not arise in the distributed network account we are suggesting. Consider, for example, bilingual orthographic

neighborhood effects (for example, see French and Ohnesorge, 1996, among many other papers on the subject). It has been shown that the recognition of a word, *W*, belonging to L1, the active language, can be significantly affected by a large orthographic neighborhood in L2, the non-active language. Now, according to the BIA+, how would this work? If the language node corresponding to L2 is not active, where did the inhibition of *W* come from? On the other hand, the authors might argue that the L2 node is partially active, due to the partial activation of the *W*'s orthographic neighbors in L2, but that its activation is below conscious threshold. Nonetheless, if they claim that this sub-conscious activation level is sufficient to have an inhibitory effect on the word in L1, then what is the use of having a separate language node whose partial activation correlates perfectly with the activation of the orthographic neighbors? Why not eliminate the language node altogether, since the real work is being done, not by a partially activated language node but, rather, by the activated and partially activated elements *at the word level*?

Another point concerns the addition of a "semantics" module. It is certainly true that a word's semantic features should be taken into account in a word identification model. For this reason, the BIA+ is, indeed, more complete than the original BIA. However, this addition comes with a price. For example, exactly *how* is semantics to be included in the model? In particular, given that theirs is a localist account, does this mean that there is one meaning per node? In this case, how are the relations between and overlap among concepts implemented? In a distributed model with learning, one can at least suggest ways in which this could be accomplished, but how do the authors propose dealing with this problem in the non-learning, localist framework in which they are working?

The final point concerns BIA+'s "task module". While it is unquestionably appropriate to consider the role of task demands, is a new module really necessary to accomplish this? Instead, different sets of elementary processing operations could be activated depending on the task requirements. In this way, task processing becomes "distributed" task processing. Task demands would then influence the operation of the system, as they should, but there would be no need for a separate "task module". A good example of a memory system with integrated task specificity is Minerva II (Hintzman, 1984), in which there is no "task module" per se, instead a different combination of processes is triggered depending on the intended goal.

Learning

Our most significant problem with the BIA+ model is its lack of learning mechanisms. This means, inevitably, that all the model's inhibitory and excitatory connections between items at the same level and between levels must be set by hand, with all the potential problems that this can entail. One of the problems the authors run into with this approach is the necessity of encoding inter-lexical homographs twice, once in the L1 module and once in the L2 module (see section 2.5 of their article). Is this reasonable from an interactive activation perspective? At the letter and feature level, a Dutch-English homograph, such as "ROOM" (meaning "cream" in English) has rigorously identical characteristics. Then, suddenly, at the word level, there are *two* distinct lexical items, one for Dutch and one for English. The only justification for this particular hand-coding is that the *authors themselves* know that "ROOM" is an interlexical homograph, and, in order for their model to produce the interference effects observed in Dutch-English bilinguals, this particular lexical item (as opposed to items like "COW") must be represented separately in the two separate language modules within their model. It strikes us that there should be a single lexical item, "ROOM", that would become active (along with the appropriate language-specific semantics) depending on the active language context, a situation that would occur if the model learned the items of the languages and the relationships among them on its own. One could imagine, for example, the statistical learning model of Christiansen, Allen and Seidenberg (1998) being extended to the case of bilingual language learning. In short, our view is that it would be far better for this type of model *to develop its own representations through learning*, as has often been proposed for the acquisition of a single language. Not only would this avoid potential problems with hand-coding but the connection strengths would be a function of directly experiencing the languages (in written form), rather than the work of programmers attempting to make the program work correctly.

As we discussed in the previous section, without learning, how do the authors propose to incorporate the semantic level? Will they hand-code all necessary semantic representations and their relationships to each other, each with varying strengths? This leads to the problem of scaling up to real languages. While it might be potentially possible for programmers to hand-code the relationships among a very small number of features, letters and lexical items of two different languages, it is much more difficult – almost certainly impossible – to do so as the number of items increases towards those required for a real language. If the program were able to learn to set its own synaptic weights, the problem of scaling-up associated with hand-coding could quite possibly be overcome.

Finally, we propose looking at the question of an integrated vs. separate lexical store. Indeed, the authors argue for integrated lexicons, but one has to wonder what, exactly, they mean by this. We will examine this question next, again in the context of learning.

Integrated vs. separate lexicons

An overwhelming body of recent experimental evidence in bilingual memory militates in favor of the idea of *unified bilingual lexicons* (Hermans, Bongaerts, de Bot and Schreuder, 1998; Costa and Caramazza, 1999; Costa, Caramazza and Sebastián-Gallès, 2000; Colomé, 2001), even though, functionally, bilingual memory gives the appearance of lexical separation. In the BIA+ this is achieved by having common feature and letter levels for the lexical items in both languages. Thereafter, at the word level, functional language separation seems for all the
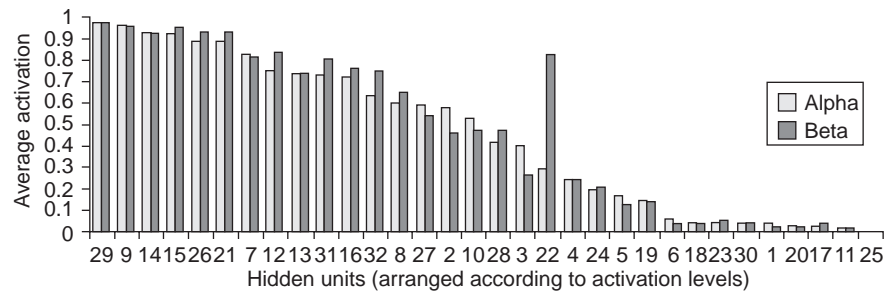
Figure 1. Overall overlap of the hidden-unit activations of all word represetations for both languages.

world to have been implemented as *physical* lexical separation. In other words, at the word level, the BIA+ seems to posit separate lexicons. We suggest that this language-based word-level separation would be unnecessary in a distributed connectionist setting.

But how could a connectionist network's highly overlapping internal representations of the items in both languages *also* produce the high degree of functional language separation that we experience as a bilingual? French (1998) used a simple recurrent connectionist network (Elman, 1990) to show how this question might be approached from the standpoint of bottom-up bilingual language learning. This very simple model could arguably serve as a starting point for a statistical, emergent approach to bilingual language learning and storage. French produced two micro-languages, Alpha and Beta, and generated random series of sentences in each language, switching from one language to the other from time to time, as bilinguals do. There were no explicit markings of sentences or language switches. Below is a sample of text presented to the network:

BOY LIFTS TOY MAN SEES PEN MAN TOUCHES BOOK GIRL PUSHES BALL WOMAN TOUCHES TOY BOY PUSHES BOOK FEMME SOULEVE STYLO FILLE PREND STYLO GAR ON TOUCHE LIVRE FEMME POUSSE BALLON FILLE SOULEVE JOUET WOMAN PUSHES PEN BOY LIFTS BALL WOMAN TAKES BOOK

After exposing the network to some 60,000 items as above, he then examined the average hidden unit representations (i.e., activation patterns) over all words in each language and compared them. The amount of overall overlap between the representations of the words in the two languages was enormous, giving the impression of very highly overlapping lexicons (see Figure 1 above).

In contrast, however, when a cluster analysis of the internal representations of the words in each language was performed, French discovered that not only were all of the parts of speech in each language clustered correctly, but also the two languages themselves had developed the exactly appropriate clusters, as can be seen in Figure 2.

The simulation was later run with two larger micro-languages containing 768 words apiece and a distributed (i.e., non-localist) encoding for the words in each language. The same high-dimensional language separation emerged spontaneously. French was also able to show inter-lingual homograph priming effects with this extremely simple model. The point is that it is at least *possible* to develop bilingual memory models along these lines.
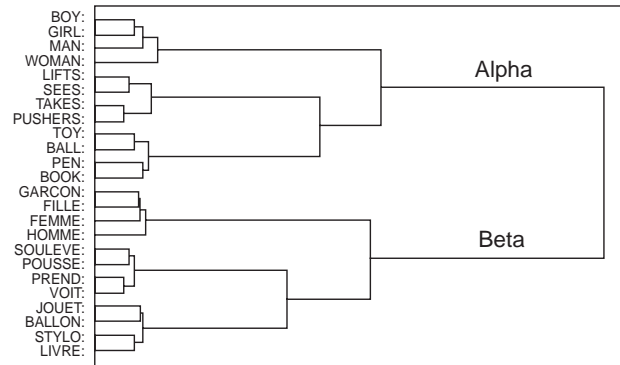

Figure 2. Cluster analysis of the hidden-unit representations for the words in both languages.

Conclusion

Dijkstra and van Heuven's BIA+ model is a well-thought out attempt to incorporate low-level mechanisms into a high-level model of bilingual memory that accounts for the extensive experimental data showing inter-lingual priming and interference effects. Just as the original Interactive Activation model paved the way for the connectionist revolution in cognitive psychology, we hope that the ground-breaking work of these authors will naturally evolve towards broader-based distributed connectionist network models and related dynamical models of bilingual memory, capable of learning and being able to incorporate both the bottom-up and the top-down processing that we know to be an integral part of bilingual language processing.

Acknowledgments

References

Chalmers, D. J., French, R. M. & Hofstadter, D. R. (1992). High-level perception, representation, and analogy: A critique of Artificial Intelligence methodology. *Journal of Experimental and Theoretical and Artificial Intelligence*, 4, 185–211.

Christiansen, M. H., Allen, J. & Seidenberg, M. S. (1998). Learning to segment speech using multiple cues: A con-

nectionist model. *Language and Cognitive Processes*, 13, 221–268.

Churchland, P. (1995). *The engine of reason, the seat of the soul*. Cambridge, MA: MIT Press.

Colomé, A. (2001). Lexical activation in bilinguals' speech production: language-specific or language-independent? *Journal of Memory and Language*, 45, 721–736.

Costa, A. & Caramazza, A. (1999). Is lexical selection in bilinguals language specific? Further evidence from Spanish-English bilinguals. *Bilingualism: Language and Cognition*, 2, 231–244.

Costa, A., Caramazza, A. & Sebastián-Gallès, N. (2000). The cognate facilitation effect: implications for model of lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1283–1296.

Dijkstra, A. (Ton) & Van Heuven, W. J. B. (1998). The BIA model and bilingual word recognition. In J. Grainger & A. M. Jacobs (eds.), *Localist connectionist approaches to human cognition*, pp. 189–225. Mahwah, NJ: Erlbaum.

Elman, J. (1990). Finding structure in time. *Cognitive Science*, 14, 179–211.

Elman, J. (1995). Language as a dynamical system. In R. Port & T. van Gelder (eds.), *Mind as motion: Explorations in the dynamics of cognition*, pp. 195–223. Cambridge, MA: MIT Press.

French, R. M. (1998). A Simple Recurrent Network model of bilingual memory. In *Proceedings of the Twentieth Annual Cognitive Science Society Conference*, pp. 368–373. Hillsdale, NJ: Lawrence Earlbaum.

French, R. M. & Ohnesorge, C. (1996). Using interlexical non-words to support an interactive-activation model of bilingual memory. *Proceedings of the Eighteenth Annual Cognitive Science Society Conference*, pp. 318–323. Hillsdale, NJ: Lawrence Erlbaum.

Hermans, D., Bongaerts, T., de Bot, K. & Schreuder, R. (1998). Producing words in a foreign language: Can speakers prevent interference from their first language? *Bilingualism: Language and Cognition*, 1, 213–230.

Hintzman, D. L. (1984). Minerva II: A simulation model of human memory. *Behavior Research Methods, Instrument & Computers*, 16, 96–101.

McClelland, J. & Rumelhart, D. (1981). An Interactive-Activation model of context effects in letter perception, part 1: An account of basic findings. *Psychological Review*, 88, 375–405.

Van Heuven, W. J. B., Dijkstra, A. (Ton) & Grainger, J. (1998). Orthographic neighbourhood effects in bilingual recognition. *Journal of Memory and Language*, 39, 458–483.

The bilingual as an adaptive system

DAVID W. GREEN
Department of Psychology, Gower Street, University College London, London WC1E 6BT, UK.
E-mail: *d.w.green@ucl.ac.uk*

Dijkstra and van Heuven lucidly summarize the important research generated by the BIA model and provide an excellent case for the BIA+ model with its critical separation of the identification system from the task/decision system. A keynote article necessarily offers a selective exposition of the authors' thinking and so my remarks are an invitation to expand. My first question concerns the scope of the BIA+ model. My remaining questions broadly address a key feature of the BIA+ model – its ability to explain performance changes as a function of task demands

Any cognitive system is a product of neural constraints. Ultimately, then, it seems to me, we need to construct a neurocognitive model of bilingualism. My first question concerns the relationship between the BIA+ (cognitive) model and neuroimaging data. I shall adopt the devil's advocate position. Consider the identification system. It includes a subsytem that explicitly represents visual word forms. Recently, Cohen et al. (2000, 2002) proposed that the posterior region of the left mid-fusiform gyrus is the visual word form area because it activates when individuals silently read briefly presented visual words or read words compared to false fonts. However, this area also activates in tasks that have nothing to do with reading words. For instance, it activates when individuals name the colours of meaningless gratings compared to saying OK to the same stimuli. It also activates when individuals press a button to indicate whether a depicted unfamiliar non-object could be twisted or poured (see Price and Devlin, submitted). Suppose there is no one cortical area corresponding to the visual word form area. Instead functional specialisation emerges from unique interactions amongst a set of regions each fulfilling a number of different functions (e.g., Mesulam, 1990). On the face of it, if that were true, it would seem to be more consistent with parallel distributed views of the visual word form system rather than the localist representations of the BIA+ model. Non-localist models dispense with the explicit representation of word form and propose that our knowledge of letter combinations arises from interactions between orthographic, semantic and phonological processing Seidenberg and McClelland, 1989; Plaut, McClelland, Seidenberg and Patterson, 1996). To what extent do Dijkstra and van Heuven see neuroimaging data as relevant to the BIA+ model? To what extent do they see their cognitive model as contributing to the wider endeavour of constructing a neurocognitive model?

My second question concerns the relationship between orthography and phonology in bilinguals. There is reasonable agreement that visual word recognition involves phonology. Indeed competition in the mapping of orthography to phonology may explain the finding that phono-logical priming effects, compared to orthographic priming effects, are smaller in bilingual speakers (e.g., Brysbaert, in press). As Dijkstra and van Heuven note, there is also reasonable evidence that the "grain-size" of processes that map orthography onto phonology differ between languages. Languages differ in the consistency of their mappings between orthography and phonology. For instance, English is less consistent than German, Italian, Greek or Spanish. In the case of a consistent orthography, phonology can be derived by mapping individual graph-emes (i.e., abstract letter identities) onto phonemes. So, for example, the "a" in Ball, Hand and Park are pronounced in the same way in German but in different ways in English. It follows that the grain size of the orthographic units required to derive a correct phonology may differ between the languages. In the case of regular languages, such as German, successive graphemes can be mapped onto phonemes. Naming time will, therefore, tend to increase with word length (strictly speaking it should increase with the number of grapheme-phoneme correspondences). In contrast, in the case of a language such as English, with an inconsistent orthography, a different grain size may be required in order to produce the correct pronunciation. For instance, individuals may need to retrieve the phonology of the word body. A word body is essentially the orthographic representation of the phonological rhyme (e.g., for Ball, the word body is "-all"). On other occasions, individuals may need to retrieve the phonology associated with the whole word.

Ziegler, Perry, Jacobs and Braun (2001) put these conjectures to the test. For German and English, Ziegler et al. were able to use exactly the same set of words (e.g., Ball, Hand, Park). They showed that word length predicted naming time for native German speakers and that body-N (the number of body neighbours of the word) predicted naming time for native English speakers. There was a 13 ms difference between three-letter and five-letter words for English speakers (pooling over body-N) and a 50 ms difference for German speakers. In contrast, pooling over word length, words with a large number of body neigh-bours were read 15 ms more quickly than those with a small number of body neighbours by English speakers but only 5 ms more quickly by German speakers.

Grain size is also important in other tasks. Rey, Ziegler and Jacobs (2000) showed that detecting a single letter (A) in a multi-unit grapheme (EA) in a word such as BEACH, is slower than detecting the same letter when it corresponds to a single grapheme as in the word PLACE. Readers may also adopt mappings at different grain-sizes in a visual search task as a function of the type of letter string presented. When individuals are required to detect a letter in a subsequently presented string of five random letters,

the reaction time function relating correct detection to the position of the target letter looks like an M sloping upwards left-to-right (Green, Rickard Liow, Tng and Zielinski, 1996). Search times are faster and the function linear when words are presented (Rickard Liow, Yap and Green, in preparation). Differences in orthographic consistency between languages raise empirical questions in the study of bilingual speakers. So here is my second question: what predictions does the BIA+ model make (given its coding scheme for sublexical orthography and phonology) about the effects of grain size for bilingual speakers? How adaptive is the system in this respect? For instance, will bilingual speakers of an L2 necessarily show the grain-size effects of native speakers of that language?

My third question concerns the role of language membership information. Dijkstra and van Heuven argue that language membership information arrives too late to affect word recognition and that its effects would, in any case, be too weak given that there is a one-to-many mapping from the language node onto lexical representations. The latter claim is plausible but the claim that language information arrives late may require qualification. Think for a moment of the role of sublexical information. Distinctive sublexical information speeds word recognition. But such information can signal language membership and so, in principle, information on language membership could become available before word recognition (see also work cited in Dijkstra and van Heuven on auditory word recognition). In consequence, it could affect response output at a very early stage. For instance, a change in language could signal a change in response and so lead individuals to respond more quickly when there is a response change and a language change compared to when there is language change but no change in response. One might also expect individuals to alter their response to the cue when it proves non-predictive (see von Studnitz and Green, 2002a for evidence compatible with these ideas).

Dijkstra and van Heuven make the strong claim that items of the two languages are always fully activated. This is a strong claim because, if I have understood correctly, it implies that even when a speaker has not used a language for years, the lexical representations of that language compete, just as strongly as they ever did, with the current language to account for the input. Experimentally, any effects of introducing pure non-target language words in a language-specific lexical decision task must be attributable to changes in the decision criterion or some aspect of the task schema. In line with the research on Dutch-English bilinguals, von Studnitz and Green (2002b) found that presenting a pure non-target German word in a list of words for English lexical decision, increased RT to an interlingual homograph such as (TAG – "day" in German) relative to an English control word, matched to the homograph's English frequency. It also increased the time required to respond to a high-frequency English word following the homograph – there was a carry-over effect. In the second half of the experiment, individuals reduced the interference effect but there was no change in the carry-over effect. One interpretation of these data is that individuals tuned a conjoint test. As Dijkstra and van Heuven

noted, such a test fires when there is evidence of a word in both languages. Such a test acts after word recognition to speed response but does not affect identification. This proposal is fully consistent with the strong separation of the identification system from the task schema proposed in the BIA+ model. However, why – if both languages are fully active – was carry-over greater in this condition compared to a condition where there was no pure German word? Von Studnitz and Green (2002b) proposed that such increased carry-over arose because there was increased activation of either the lexical representations of the non-target language or the units coding for language membership (i.e., a language tag) and that this activation fed into the units mediating the "no" response. Because of competition to control output, residual activation in these units increased the time required to respond "yes" on the following trial.

Dijkstra and van Heuven rule out the possibility of increased activation of lexical units. But is the second possibility allowed? Couldn't stimulus input increase the activation of the units coding for language membership? Whereas there is a one-to-many mapping of these units onto lexical representations, there can be a one-to-one mapping of these units onto a response (e.g., if English, press "yes"). If the activation of these units is also invariant, Dijkstra and van Heuven can explain the carry-over effect in another way, for example, as a result of a temporary change in the decision criterion for real high-frequency English words though such an explanation is perhaps rather subtle.

Finally, the arena of individual differences is an important one in bilingual research. A bilingual mental system arises under specific contexts of acquisition and use. Typically, localist models are silent on such broadly developmental matters. Could Dijkstra and van Heuven say a little more about individual differences in visual word recognition in bilinguals in the light of the BIA+ model?

References

Brysbaert, M. (in press). Bilingual visual word recognition: Evidence from masked phonological priming. In S. Kinoshita & S. J. Lupker (eds.), *Masked priming: State-of-the-art*. Hove: Psychology Press.

Cohen, L., Dehaene, S., Naccache, L., Lehericy, S., Dehaene-Lambertz, G., Henaff, M. & Michel, F. (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.

Cohen, L., Lehericy, S., Chochon, F., Lemer, C., Rivard, S. & Dehaene, S. (2002). Language-specific tuning of visual cortex? Functional properties of the visual word form area. *Brain*, 125, 1054–1069.

Green, D. W., Rickard Liow, S., Tng, S. K. & Zielinski, S. (1996). Are visual search procedures adapted to the nature of the script? *British Journal of Psychology*, 87, 311–326.

Marshall, J. C. & Newcombe, F. (1973). Patterns of paralexia: A psycholinguistic approach. *Journal of Psycholinguistic Research*, 2, 175–199.

Mesulam, M. M. (1990). Large scale neurocognitive networks and distributed processing for attention, language and memory. *Annals of Neurology*, 28, 597–613

Plaut, D., McClelland, J., Seidenberg, M. & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56–115.

Price, C. J. & Devlin, J. T. (submitted). The myth of the visual word form area.

Rey, A., Ziegler, J. C. & Jacobs, A. M. (2000). Graphemes as perceptual reading units. *Cognition*, 75, B1–B12.

Rickard Liow, S. J., Yap, M. J-M. & Green, D. W. (in preparation). Word identification and letter search: Effects of lexicality, legality, but not frequency.

Seidenberg, M. & McClelland, J. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96, 523–568.

von Studnitz, R. E. & Green, D. W. (2002a). The cost of switching language in a semantic categorization task. *Bilingualism: Language and Cognition*, 5, 241–251.

von Studnitz, R. E. & Green, D. W. (2002b). Interlingual homograph interference in German-English bilinguals: Its modulation and locus of control. *Bilingualism: Language and Cognition*, 5, 1–23.

Ziegler, J. C., Perry, C., Jacobs, A. M. & Braun, M. (2001). Identical words are read differently in different languages. *Psychological Science*, 12, 379–384.

Bilingual word recognition beyond orthography: On meaning, linguistic context and individual differences

JANET G. VAN HELL
University of Nijmegen, Department of Special Education,
PO Box 9104, 6500 HE Nijmegen, The Netherlands.
E-mail: *J.vanHell@ped.kun.nl*

Central questions in psycholinguistic studies on bilingualism are how bilinguals access words in their two languages, and how they control their language systems and solve the problem of cross-language competition. In their excellent paper "The architecture of the bilingual word recognition system: From identification to decision", Dijkstra and Van Heuven expound their BIA+ model on bilingual word recognition. BIA+ builds on its predecessor BIA, one of the first connectionist models on bilingual word recognition. BIA+ preserves one of BIA's crucial assumptions, namely that the bilingual lexicon is integrated across languages and is accessed in a language non-selective way, an assumption that is supported in many empirical studies and that is now widely accepted in the bilingual literature. Compared to the original BIA model, the BIA+ architecture is further developed (in fact, much more so than the subtle 'plus' denotes). BIA+ now includes orthographic, as well as phonological and semantic representations in the word identification system, and a distinction is made between a word identification system and a task/decision system. This latter extension resembles the language task schemas in Green's (1998) Inhibitory Control model. Dijkstra and Van Heuven also distinguish between effects of linguistic and non-linguistic context on performance: linguistic context effects, that arise from lexical, syntactic and semantic sources, are assumed to affect the activity in the word identification system, whereas non-linguistic effects, that can arise from instruction, task demands or participant expectancies, are assumed to affect the task/decision system.

Dijkstra and Van Heuven's BIA+ model is a comprehensive model of bilingual word recognition and opens new ways of approaching word recognition in bilinguals and the control mechanisms involved. A compelling perspective of the model is that a bilingual is not conceived of as a computational system that operates in isolation and is intrinsically context-free, but as a system that acts in context and continuously interacts with and adapts to linguistic and non-linguistic contextual factors. The latter view is also a central assumption in the dynamical system perspective on cognition and language (e.g., Elman, 1995; Van Gelder, 1998). Though it may not always be easy to disentangle linguistic from non-linguistic information in a language situation (cf. the effects of stimulus list composition), Dijkstra and Van Heuven's description of specific contextual factors and their influence on the word identification process and task/decision system inspires the formulation of hypotheses to be tested in future studies.

Specifically, how words are recognized in natural sentence and discourse contexts is an important goal in future research and may form a critical test of the validity and plausibility of bilingual word recognition models, including BIA+.

In this commentary, I mainly focus on the word identification system and substantiate Dijkstra and Van Heuven's main points by bringing in and elaborate on some central issues in dynamical and connectionist approaches to cognition and language: time, flexibility, context and idiosyncrasy. Specifically, I will focus on bilingual word identification in light of temporality, the flexibility of word meanings and cross-linguistic code overlap, contextual information and grammatical class, and individual differences among bilinguals.

The word identification system in the BIA+ model consists of orthographic, phonological and semantic representations, which implies that word recognition can be affected by spelling, sound and meaning, an assumption also made by models on monolingual word recognition in adults (e.g., Van Orden and Goldinger, 1994) and children (e.g., Bosman and Van Orden, 1997; Bosman and Van Hell, in press). Dijkstra and Van Heuven propose that the connections between these three codes are interactive and bi-directional, and that the three codes interact (resonate) over time. At a general level, the resonance assumption has important implications with respect to theoretical perspectives on how language behavior unfolds in time, and differs in fundamental ways from the perspective adopted in feedforward networks (see, e.g., Elman, 1995 for more details). Moreover, the postulation of bi-directional connections has important consequences for the activation flow within the word identification system: activation flows from orthographic codes to phonological and semantic codes (feedforward activation), and also feedback from semantic and phonological codes to orthography (feedback activation). Thus, following initial activation, recurrent feedback dynamics begin among all these three codes. The bi-directional flow of activation among the three codes and the role of feedback activation in visual word recognition is under lively debate in the monolingual literature, for example in relation to (in)consistencies in the mappings between phonology and orthography (Stone, Vanhoy and Van Orden, 1997) and semantics and orthography (e.g., Pecher, 2001). Though the role of feedback activation remains largely implicit in Dijkstra and Van Heuven's model, given that BIA+'s word identification system is highly interconnected across orthographic, phonological

and semantic codes and across languages, feedback activation from phonology to orthography and from semantics to orthography will also affect bilingual word recognition (cf. Tokowicz, Kroll, De Groot and Van Hell, 2002).

The addition of semantic representations is an important enrichment in BIA+, and enhances its explanatory power considerably. Dijkstra and Van Heuven discuss semantic representations mainly in the light of their findings on interlingual homographs (like *room* meaning 'cream' in Dutch) and cognates (like Dutch-English *film*), and assume that interlingual homographs do not share meaning and cognates do. Recent studies suggest, however, that cross-linguistic semantic overlap ranges across a continuum (e.g., De Groot, Dannenburg and Van Hell, 1994; Van Hell and De Groot, 1998) and that even at the endpoint of the scale, featuring the cognates, semantic overlap is graded. More specifically, Van Hell and De Groot (1998) asked Dutch-English bilinguals to associate twice to the same stimuli on different occasions, once in the language of the stimuli (within-language) and once in the other language (between-language). It appeared that the within- and between-language associations for concrete cognates were more often translations of one another than those for abstract cognates, and retrieving an associate was easier for concrete cognates than for abstract cognates. This suggests that cognates can differ in their degree of semantic code overlap.

Though Dijkstra and Van Heuven do not explicitly describe the representational format the semantic information is coded in, the assumption of cross-linguistic code overlap and co-activation within the semantic system would fit in with current ideas on the nature of word meanings, and with the basic architecture of the distributed feature model on semantic representation in bilingual memory. The classical idea that words have a fixed meaning and that the meaning of a word can be captured by defining a set of individually necessary and jointly sufficient features is called into question by an increasing number of scientists, on theoretical or on empirical grounds. Opponent theorists objected that it is very difficult to construct lists of defining features for words, and argued, for example, that there is no principled way to distinguish essential features from non-essential features. They brought up white tigers that violate the 'being-striped' feature but that are still considered tigers (Armstrong, Gleitman and Gleitman, 1983), or concepts that meet all features, but still would not be adorned with the corresponding label (e.g., "Is Tarzan a bachelor?", see Lakoff, 1987). Moreover, Labov (1973) obtained compelling evidence that concept boundaries are graded and vary with the context of use. Word meanings can thus be considered "slippery customers" (Labov, 1973, p. 341): rather than consisting of a set of defining features, word meanings are vague, have fuzzy boundaries and depend on the context.

The idea that word meanings are slippery customers is central in recent views on conceptual representation in memory (e.g., Barsalou, 1993; Hummel and Holyoak, 1997). Rather than assuming that word meanings are stable structures that are stored in long-term memory and retrieved when needed, it is proposed that word meanings

are dynamically constructed in working memory, and can vary widely in the conceptual features that are incorporated. This notion predicts that people conceptualize the same word differently on different occasions. This is corroborated in the above-mentioned word association study (Van Hell and De Groot, 1998): it appeared that associative responses were more often not repeated than repeated, even in the condition in which participants performed the same task twice in their native language. Furthermore, associations to abstract words, noncognates and verbs were less often repeated within- and across-languages than those to concrete words, cognates and nouns. We explained these findings in terms of the distributed feature model, in which the meaning of words is represented in terms of distributed features. Word meaning is conceived of as graded and flexible, and the similarity of the semantic representations of two words is a function of the conceptual feature overlap of the words. It is further proposed that cross-linguistic conceptual feature overlap ranges across a continuum, and that the overlap of some translation equivalents, e.g., concrete nouns and cognates, is higher than that of others, e.g., abstract nouns and noncognates (see, e.g., De Groot, 1992; Kroll and De Groot, 1997; Van Hell, 1998; Van Hell and De Groot, 1998 for more extensive discussions). Cross-linguistic overlap may thus range across a continuum in the orthographic and phonological codes (e.g., Dijkstra, Grainger and Van Heuven, 1999; Van Hell and Dijkstra, in press), as well as semantic codes. Moreover, the set of features that is activated is not a static entity, and may depend on factors like the context in which the word is embedded.

As already mentioned above, its emphasis on the explicit role of contextual factors in bilingual language performance is a cogent quality of BIA+. In the recent past, considerable progress has been made in understanding bilingual recognition and production of isolated words. Surprisingly few studies focused on how word recognition or production operates in context, in particular sentence or discourse context. In their discussion of linguistic context effects, Dijkstra and Van Heuven propose that the word identification system interacts with the sentence parsing system, and that the recognition of words in sentence context is sensitive to syntactic and semantic information from both languages: word recognition and syntactic and semantic effects are language-nonselective. This proposal opens up important questions for future research, one of them being the question Dijkstra and Van Heuven discuss in some detail, namely how sentence constraints influence the generation of featural restrictions for upcoming words and how they thus delineate the activation of lexical and conceptual elements within and across languages.

A related question pertains to the temporal aspects of sentence processing and word identification processes in L2 and L1. To describe the delayed activation of L2 phonological and semantic codes relative to the L1 codes in the word identification system, Dijkstra and Van Heuven introduce a temporal delay mechanism. Does this mechanism also affect the activation of syntactic and semantic information during sentence processing? What are the consequences for cross-linguistic effects in L1 and L2 sentence

processing? Furthermore, Dijkstra and Van Heuven's model opens up the question of how the sentence parser operates in bilinguals. Do bilinguals, and in particular those who learned their L2 at a time when their L1 processing system was already well-developed, like the late bilinguals that were tested in the experiments that inspired BIA+, use their L1 parsing routines when reading L2 sentences, or do they use L2 parsing routines? This issue points to the importance of control in solving cross-language competition of the sentence parsing routines. An early and partial answer to such questions is provided by studies in which L2 speakers are compared with monolingual speakers of the same language (see Kroll and Dussias, in press, for a comprehensive review). The available evidence shows that in processing semantic incongruities, differences between the two speakers were quantitative only and arose because semantic processes were slowed down in L2 speakers (e.g., Weber-Fox and Neville, 1996). Evidence from studies investigating syntactic parsing is less conclusive, however: L2 speakers have been found, for example, to behave like monolingual speakers, to behave qualitatively differently, or to transfer L1 information when processing in L2. What these studies at least suggest is that language-specific syntactic properties (e.g., in verb-argument structure) may challenge the universality and language-independence of the processing routines that are used to read sentences in L1, in L2 or in both languages.

A discussion of the role of linguistic context effects in word recognition draws attention to the different grammatical functions of words in a sentence. Behavioral research on bilingual word recognition (and production) largely ignored the grammatical characteristics of words and used words from one grammatical class only, nouns. This implies that theoretical models of word recognition, including BIA+, in fact describe noun recognition. Monolingual, cross-linguistic and bilingual studies have found, however, that lexical processing is different for nouns and verbs, and suggested that verbs (often denoting relational concepts) have a greater breadth of meaning than concrete nouns (denoting referential concepts; Gentner, 1981), and that the meanings of verb translations are less similar across languages than those of concrete nouns (Van Hell and De Groot, 1998). Furthermore, lexical processing of open class words (content words: nouns, verbs, adjectives) differs from closed class words (function words, including articles, conjunctions and prepositions; e.g., Pulvermüller, 1999), and context effects are not necessarily the same for these two groups of words (Van Petten and Kutas, 1991). Such effects of grammatical class may constrain existing models of word recognition, including bilingual word recognition, and may affect theories of word recognition and syntactic and semantic effects in sentence processing. A profound understanding of how words are recognized in context and how the word identification system interacts with syntactic and semantic information thus not only asks for research on word recognition in context, but also for further research on isolated word recognition in which the grammatical class of words is varied.

A final point I would like to bring up pertains to individual differences among bilinguals, including different levels of L2 fluency. Though the primary focus of BIA+ is not to explain how word recognition is modulated by individual differences, Dijkstra and Van Heuven touch on some mechanisms that may account for differences between less and more proficient bilinguals: differences in the resting level activation of the orthographic representations in the word identification system or (small) variations in parameters in the decision process. These mechanisms may well explain differences in word recognition in speakers with different fluency levels (e.g., Schulpen, Dijkstra and Schriefers, cited in Dijkstra and Van Heuven's article; Van Hell and Dijkstra, in press), it remains to be seen whether they can also describe differences between more or less proficient L2 speakers processing language beyond the isolated word, as observed in sentence processing studies. For example, in an ERP study using L2 speakers with different levels of exposure to their L2 (and who also differed in proficiency), Weber-Fox and Neville (1996) found substantial amplitude and distribution differences between bilinguals with different levels of L2 exposure in syntactic processing (although brain systems mediating semantic processing of sentences appeared relatively robust against different levels of L2 exposure). In addition to differences in L2 exposure and proficiency, several recent studies suggest that individual differences between bilinguals in L2 learning rate (Osterhout, McLaughlin, Inoue and Loveless, 2001) and reading span (Kroll, Michael, Tokowicz and Dufour, 2002) may modulate cross-language processing. In their paper, Dijkstra and Van Heuven point out that one of the aspects that should be considered in the future is how BIA+ would develop over time and during learning. Though I fully concur in the importance of modeling the continuing process of language learning and development in novice and fluent bilinguals, I would add to the agenda a focus on individual differences among bilinguals.

From the discussion above it is evident that Dijkstra and Van Heuven's inspiring model opens new approaches to the structure and dynamics of bilingual word recognition beyond the isolated word. A further testing of the ideas outlined in their paper will enhance our understanding of bilingual language processing.

References

Armstrong, S. L., Gleitman, L. R. & Gleitman, H. (1983). What some concepts might not be. *Cognition*, 13, 263–308.

Barsalou, L. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. F. Collins, S. E. Gathercole, M. A. Conway & P. E. Morris (eds.), *Theories of memory*, pp. 29–101. Hove: Lawrence Erlbaum.

Bosman, A. M. T. & Van Hell, J. G. (in press). Orthography, phonology, and semantics: Concerted action in word perception. In L. Verhoeven, C. Elbro & P. Reitsma (eds.), *Precursors of functional literacy*. Amsterdam & Philadelphia: John Benjamins.

Bosman, A. M. T. & Van Orden, G. C. (1997). Why spelling is more difficult than reading. In: C. A. Perfetti, L. Rieben & M. Fayol (eds.), *Learning to spell: Research, theory, and practice across languages*, pp. 173–194. Hillsdale, NJ: Lawrence Erlbaum.

De Groot, A. M. B. (1992). Bilingual lexical representations: A closer look at the conceptual representations. In R. Frost & L. Katz (eds.), *Orthography, phonology, morphology, and meaning*, pp. 389–412. Amsterdam: Elsevier.

De Groot, A. M. B., Dannenburg, L. & Van Hell, J. G. (1994). Forward and backward word translation by bilinguals. *Journal of Memory and Language*, 33, 600–629.

Dijkstra, A., Grainger, J. & Van Heuven, W. J. B. (1999). Recognition of cognates and interlingual homographs: The neglected role of phonology. *Journal of Memory and Language*, 41, 496–518.

Elman, J. L. (1995). Language as a dynamical system. In R. F. Port & T. Van Gelder (eds.), *Mind as motion: Explorations in the dynamics of cognition*, pp. 195–225. Cambridge, MA: MIT Press.

Gentner, D. (1981). Some interesting differences between verbs and nouns. *Cognition and Brain Theory*, 4, 161–178.

Green, D. W. (1998). Mental control and the bilingual lexico-semantic system. *Bilingualism: Language and Cognition*, 1, 67–81.

Hummel, J. E. & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104, 427–466.

Kroll, J. F. & De Groot, A. M. B. (1997). Lexical and conceptual memory in the bilingual: Mapping form to meaning in two languages. In A. M. B. De Groot & J. F. Kroll (eds.), *Tutorials in bilingualism: Psycholinguistic perspectives*, pp. 169–199. Hillsdale, NJ: Lawrence Erlbaum.

Kroll, J. F. & Dussias, P. E. (in press). The comprehension of words and sentences in two languages. In T. Bhatia & W. Ritchie (eds.), *Handbook of bilingualism*. Cambridge, MA: Blackwell Publishers.

Kroll, J. F. Michael, E., Tokowicz, N. & Dufour, R. (2002). The development of lexical fluency in a second language. *Second Language Research*, 18, 137–171.

Labov, W. (1973). The boundaries of words and their meanings. In C.-J. N. Baily & R. W. Shuy (eds.), *New ways of analyzing variation in English*, pp. 340–373. Washington, DC: Georgetown University Press.

Lakoff, G. (1987). *Women, fire, and dangerous things*. London: The University of Chicago Press.

Osterhout, L., McLaughlin, J., Inoue, K. & Loveless, J. (2001). *Brain activity in L2 learning*. Paper presented at the Annual Meeting of the Cognitive Neuroscience Society, March 25–27, New York.

Pecher, D. (2001). Perception is a two-way junction: Feedback semantics in word recognition. *Psychonomic Bulletin & Review*, 8, 545–551.

Pulvermüller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, 22, 253–279.

Stone, G. O., Vanhoy, M. & Van Orden, G. C. (1997). Perception is a two-way street: Feedforward and feedback phonology in visual word recognition. *Journal of Memory and Language*, 36, 337–359.

Tokowicz, N., Kroll, J. F., De Groot, A. M. B. & Van Hell, J. G. (2002). Number of translation norms for Dutch-English translation pairs: A new tool for examining language production. *Behavior Research Methods, Instruments, and Computers*, 34, 435–451.

Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21, 615–628.

Van Hell, J. G. (1998). Cross-language processing and bilingual memory organization. Ph.D. dissertation, University of Amsterdam.

Van Hell, J. G. & De Groot, A. M. B. (1998). Conceptual representation in bilingual memory: Effects of concreteness and cognate status in word association. *Bilingualism: Language and Cognition*, 1, 193–211.

Van Hell, J. G. & Dijkstra, A. (in press). Foreign language knowledge can influence native language performance in exclusively native contexts. *Psychonomic Bulletin and Review*.

Van Orden, G. C. & Goldinger, S. D. (1994). Interdependence of form and function in cognitive systems explains perception of printed words. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 1269–1291.

Van Petten, C. & Kutas, M. (1991). Influences of semantic and syntactic context on open- and closed-class words. *Memory and Cognition*, 19, 95–112.

Weber-Fox, C. M. & Neville, H. J. (1996). Maturational constraints on functional specializations for language processing: ERP and behavioral evidence in bilingual speakers. *Journal of Cognitive Neuroscience*, 8, 231–256.

Bilingualism is in dire need of formal models

PING LI
Department of Psychology, University of Richmond, Richmond, VA 23173, USA.
E-mail: *pli@richmond.edu*

Each year the Cognitive Science Society honors David Rumelhart by awarding the Rumelhart Prize to an outstanding cognitive scientist whose research makes a significant contribution to the formal analysis of human cognition. Formal models of language, including those of Rumelhart and his associates, are well known to psycholinguists in the monolingual context. The formalism of language in the bilingual context, however, is lamentable: to this date, there are only a handful of models (connectionist or otherwise) that are designed specifically to account for bilingual language processing and acquisition (cf. Li and Farkas, 2002). BIA is one of them. BIA+ is now another.

Bilingualism research is faced with a dilemma today. While the field lacks formal models like BIA and its variant, researchers must deal with a large number of crucial variables or factors that are believed to influence bilingual language processing, a number that starts to exceed our short-term memory buffer. These factors, as can be discerned from the keynote article, include (minimally) bilingual proficiency, history of learning, time of L2 learning, task demand in experimentation, type of words in the lexicon being tested, similarity structure of the two languages and bilingual speech mode. It is difficult to make sense of a bilingual study that does not control for these variables – indeed, many studies either ignore them in experimentation or fail to take them into account in interpretation (see discussions in Grosjean, 1998). Now the question is, how can one control for all of these variables in any study? Formalism gets us out of the dilemma with a handy control of parameters in one sweep.

The extension of BIA to BIA+ is one example of how we can explain bilingualism more precisely by considering more variables and adding newer mechanisms to the model. For example, Dijkstra and van Heuven point out that the explanatory power of BIA has been limited by its lexical representations, its ability to handle context effects and its lack of an implemented task structure. BIA+ goes beyond these limitations by incorporating phonological representations, linguistic and non-linguistic context effects, and a task-decision component.

But even this is not enough. As Dijkstra and van Heuven acknowledge, there are many other aspects that need to be considered for bilingual word recognition. One future direction they mention is how the model would develop over time in learning. In this regard, I argue that the study of language representation in adults should ultimately be connected to research in developmental bilingualism (to echo a similar point made by Bialystok, 2001). Formal models of bilingualism may take an initiative in making this connection. Such a connection can potentially reconcile conflicting results in empirical research. For example, a central argument of BIA and BIA+ is that on-line access to the bilingual lexicon is language-independent,

in contrast to the language-selective view of bilingual lexical access. However, it may be the case that developmentally, both language-independent access and language-selective access are possible, depending on the level of proficiency in bilingual language development. For highly proficient bilinguals, orthographic overlap in the bilingual lexicon leads to inhibitory effects in the access of words in the target language (Bijeljac-Babic, Biardeau and Grainger, 1997), providing evidence for language-independent access to the bilingual lexicon. For the beginning bilingual learners, however, such inhibitory effects may be much weaker or nonexistent, as shown in Bijeljac-Babic et al.'s study, in which bilingual proficiency was manipulated, thus providing evidence for language-selective access. This situation is analogous to lexical ambiguity processing in the monolingual context, where exhaustive access or selective access of multiple meanings of an ambiguous word depends on factors such as strength of context bias, frequency of each of the multiple meanings, and density of the competing meanings. Although BIA (and BIA+) is able to simulate this type of effects in bilingualism, for the most part, the model remains a "proficient bilingual model" rather than a "developmental bilingual model".

Cognitive theories of language are notoriously dichotomous. Yet language use is a complex human behavior that lies on a continuum of multiple dimensions, and this complexity is reflected even more strongly in bilingual language processing and bilingualism in development. Thus, the very phenomenon of bilingualism requires that our interpretations take into account diverse perspectives and be non-dichotomous in nature (cf. Grosjean, 1998). Formal models of bilingualism can facilitate this process, as demonstrated by BIA and BIA+. In addition, contrary to Dijkstra and van Heuven's conclusion that the implementation of the model must wait for the accumulation of empirical knowledge, I would argue that modeling allows us to consider crucial factors more effectively, and that it can generate predictions that aren't yet available in empirical research but can be tested in empirical studies.

References

Bialystok, E. (2001). *Bilingualism in development: Language, literacy, and cognition*. Cambridge: Cambridge University Press.

Bijeljac-Babic, R., Biardeau, A. & Grainger, J. (1997). Masked orthographic priming in bilingual word recognition. *Memory & Cognition*, 25, 447–457.

Grosjean, F. (1998). Studying bilinguals: Methodological and conceptual issues. *Bilingualism: Language and Cognition*, 1, 131–149.

Li, P. & Farkas, I. (2002). A self-organizing connectionist model of bilingual processing. In R. Heredia & J. Altarriba (eds.), *Bilingual sentence processing*, pp. 59–85. Amsterdam: Elsevier Science Publisher.

How do bilinguals control their use of languages?

ARDI ROELOFS
*Max Planck Institute for Psycholinguistics, P.O. Box 310,
6500 AH Nijmegen, The Netherlands.*
E-mail: *ardi@mpi.nl*

Dijkstra and Van Heuven sketch the BIA+ model for visual word processing in bilinguals. BIA+ differs in a number of respects from its predecessor, BIA, the leading implemented model of bilingual visual word recognition. Notably, BIA+ contains a new processing component that deals with task demands. BIA+ has not been computationally implemented yet and design decisions still need to be taken. In this commentary, I outline a proposal for modeling the control of tasks in BIA+.

The issue of task demands and executive control is clearly an important one. Conversations are driven by willed goals – speakers and writers try to achieve communicative intentions, and listeners and readers try to recover these intentions. Bilinguals have more than one language available for use. Thus, action goals such as addressing an interlocutor in one language need to be protected against the inadvertent use of the other language. As concerns bilingual reading, Dijkstra and Van Heuven argue that word recognition is not controlled – a letter string activates all compatible words regardless of their language. However, activated words may be responded to in multiple ways, for example, they may be read aloud, translated or be subjected to a language or lexical decision. Dijkstra and Van Heuven review evidence suggesting that the values of control parameters like decision thresholds may vary depending on the language and task situation. Clearly, what task to perform and responses to select is under the control of a language user.

Before the cognitive revolution in psychology in the middle of last century, associationist and behaviorist theories accounted for task performance and response selection by postulating associations between stimuli and responses. However, if all responses were determined entirely by stimulus-response associations, readers would not be free to choose which response to make to a letter string, because the strongest association (e.g., oral reading) would control the response. When readers are asked to perform one particular task (e.g., language decision) rather than another (e.g., oral reading), they are able to do it, usually without (much) practice. In some way or another, goals can be set to control responding.

On one prominent view, advocated by Norman and Shallice (1986) for cognition in general and adopted by Green (1998) for the control of languages in bilinguals, executive control is achieved by associatively *biasing* activation levels of responses. Goals are "internal stimuli" that favor certain response types. For example, the implemented models of Cohen, Dunbar and McClelland (1990) and Gilbert and Shallice (2002) achieve goal-directed responding by task nodes connected to response pathways in a lexical network. The task nodes selectively activate one pathway (e.g., for language decision) rather than another (e.g., for oral reading).

On another prominent view, advocated by Anderson (1983) among others and implemented in WEAVER++ (Roelofs, 1992, 1997; Levelt, Roelofs and Meyer, 1999), responding is controlled through *explicit* reference to goals. WEAVER++ combines a lexical network with condition-action production rules that determine what is done with the activated lexical information depending on the task. When a goal is placed in working memory, the attention of the system is focussed on those production rules that include the goal among their conditions. Production rules provide flexibility in responding by allowing tasks to be specified through a combination of response types and variables such as "read the stimulus aloud", "translate the stimulus from language x into y", "decide whether the stimulus is from language x", whereby x and y can take on the values "Dutch" and "English" in Dutch-English bilinguals.

Elsewhere, I have made a case for goal-referenced control of language use (Roelofs, in press). The control of language use has in its simplest form perhaps been most intensively studied by employing the color-word Stroop task (Stroop, 1935) and analogs of it. The basic variant of the Stroop task asks for naming the ink color of color words. Participants are slower and make more errors in naming the ink color of an incongruent color word (e.g., the word RED in blue ink) than in naming the color of Xs, which shows the influence of executive control: the ink colors are named rather than the color words read at the cost of slower responding and more errors.

It has been shown that WEAVER++ successfully simulates sixteen classic data sets on Stroop, mostly taken from the review by MacLeod (1991), including incongruency, congruency, reverse Stroop, response set, semantic gradient, time course, stimulus, spatial, multiple task, manual, bilingual, training, age, and pathological effects (Roelofs, in press). With only three free parameters taking two values each to accommodate task differences (color naming, picture naming, word reading, manual responding), the model accounts for ninety six percent of the variance of the sixteen studies (two hundred fifty data points). Moreover, WEAVER++ successfully simulates the human brain's blood flow response during Stroop task performance in neuroimaging studies, in particular, the fMRI BOLD response in the anterior cingulate cortex, one of the classic brain areas involved with executive control (Roelofs & Hagoort, 2002). It appears that existing associationist models of control cannot account for critical aspects of the data, whereas WEAVER++ can.

Performing a bilingual color naming Stroop task by bilinguals not only calls for naming the ink color rather than reading the word, but it also requires selecting the response from the target language. For example, when a Dutch-English bilingual has to use English in naming the

ink color of the Dutch word ROOD (red) in blue ink, correct responding not only requires preventing the selection of the response ''rood'' but also requires preventing the selection of the Dutch translation equivalent of the English target response ''blue'', Dutch ''blauw''. Stroop interference occurs between the languages of bilinguals (e.g., reading Dutch ROOD interferes with the production of ''blue'' in response to ROOD in blue ink), but the effect is not as great as that within either one of the languages.

Language-specific responding may be achieved by marking the words in memory for language, as done in BIA and WEAVER++ and proposed for BIA+, and specifying the target language in the production rules for response selection, as done in WEAVER++. Simulations have revealed that WEAVER++ accounts for the finding that interference is greater within than between languages (Roelofs, in press). As Dijkstra and Van Heuven argue for BIA+, critical is how the activated words from the non-target language are *used*. In performing a bilingual Stroop task, the words in the non-target language are not considered for selection but still compete indirectly by activating competitor responses in the target language via their shared meaning (e.g., Dutch ''rood'' activates English ''red'', which competes with English ''blue'' in naming the ink color of the Dutch word ROOD in blue). The indirect competition by words from the non-target language explains why the interference is reduced between compared to within languages.

Stroop interference between languages supports BIA+'s and WEAVER++'s assumption that words of the non-target language are processed up to the semantic level in reading. This assumption agrees with the results of a wide range of chronometric studies on bilingual word recognition, reviewed by Dijkstra and Van Heuven, but it is at odds with a recent electrophysiological and fMRI study by Rodriguez-Fornells, Rotte, Heinze, Nösselt and Münte (2002). In this study, bilingual Spanish-Catalan and monolingual Spanish readers were instructed to press a button when reading words in Spanish and to ignore pseudowords and words in Catalan. The evoked electrical brain potentials of both the bilinguals and the monolinguals were sensitive to the frequency of the words in Spanish but not in Catalan. Furthermore, the brain's blood flow response to Spanish words was greater in the posterior inferior frontal cortex and the planum temporale for the bilinguals than for the monolinguals. These areas have been shown to be involved in phonological processing by previous research. According to Rodriguez-Fornells et al., these results suggest that bilingual readers may prevent activation of the words from the non-target language by selectively accessing the lexicon of the target language through application of language-specific sublexical grapheme-to-phoneme correspondence rules and blocking direct visual lexical access for both languages.

However, if bilingual readers can control reading routes, the simplest way to prevent between-language interference in Stroop color naming would be to block all reading routes. But the Stroop findings suggest that bilingual readers are unable to accomplish this. The absence of an effect on brain potentials and the difference in brain activation between monolinguals and bilinguals does not imply, however, that bilingual word recognition is controlled. The bilingual Stroop findings indicate that between-language effects on the use of activated lexical information are reduced compared to within-language effects. Reduced effects for words from the non-target language may lead to an absence of an effect of non-target language words on brain potentials. Furthermore, the availability of two languages rather than one may lead to more extensive processing before responding in bilinguals compared to monolinguals, which may explain the fMRI data.

To conclude, the burgeoning literature on bilingual word recognition is leading theorists to think more and more about issues of control. In modeling executive control, an associationist approach is not the only way to go, however. In developing the task component of BIA+, goal-referenced control warrants serious consideration.

References

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Cohen, J., Dunbar, K. & McClelland, J. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, 97, 332–361.

Gilbert, S. M. & Shallice, T. (2002). Task switching: A PDP model. *Cognitive Psychology*, 44, 297–337.

Green, D. W. (1998). Mental control of the bilingual lexico-semantic system. *Bilingualism: Language and Cognition*, 1, 67–81.

Levelt, W. J. M., Roelofs, A. & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, 22, 1–38.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, 109, 163–203.

Norman, D. A. & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwarts & D. Shapiro (eds.), *Consciousness and self-regulation: Advances in research and theory* (vol. 4), pp. 1–18. New York: Plenum Press.

Rodriguez-Fornells, A., Rotte, M., Heinze, H.-J., Nösselt, T. & Münte, T. (2002). Brain potential and functional MRI evidence for how to handle two languages with one brain. *Nature*, 415, 1026–1029.

Roelofs, A. (1992). A spreading-activation theory of lemma retrieval in speaking. *Cognition*, 42, 107142.

Roelofs, A. (1997). The WEAVER model of word-form encoding in speech production. *Cognition*, 64, 249–284.

Roelofs, A. (in press). Goal-referenced selection of verbal action: Modeling attentional control in the Stroop task. *Psychological Review*.

Roelofs, A. & Hagoort, P. (2002). Control of language use: Cognitive modeling the hemodynamics of Stroop task performance. *Cognitive Brain Research*, 15, 85–97.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18, 643–662.

Theories that develop

MICHAEL S. C. THOMAS
School of Psychology, Birkbeck College, University of London, Malet St., London WC1E 7HX, UK.
E-mail: *m.thomas@psychology.bbk.ac.uk,*
m.thomas@ich.ucl.ac.uk

The target article represents a significant advance in the level of sophistication applied to models of bilingual word recognition, and Dijkstra and van Heuven are to be congratulated on this endeavour. Bearing in mind the success of the (computational) BIA model in capturing detailed patterns of experimental data, I look forward to future simulation results from the BIA+ when the proposals of this new framework are implemented. It is an essential step to draw a distinction between recognition systems and the decision mechanisms that drive responses, and the authors have provided a novel way of apportioning empirical evidence of context effects in bilingual word recognition across this divide. Given the explanatory weight now being placed on decision mechanisms rather than the word recognition system itself, perhaps indeed it is now time to make some simplifying assumptions about the recognition system and start building detailed computational models of the decision component of the system. Implementation will provide the clarity of theorisation and evaluation of theory viability that have been the hallmark of the BIA model thus far.

However, in this commentary, I want to focus on one particular avenue of theory development which is as yet under-explored in the BIA+ framework. This is how the cognitive structures of the bilingual's word recognition system are acquired and maintained. Dijkstra and van Heuven comment "there are many other aspects of bilingual word recognition that should be considered in the future (for instance, how the model would develop over time and during learning)" (p. 181). It is quite right, of course, to take one step at a time. It is enough work trying to explain the mass of bilingual word recognition data via a static model of the adult system, and the authors have outlined a proposal with great potential on that front. However, many of the interesting issues that characterise the bilingual language system are *dynamic* ones (Thomas & van Heuven, in press).

For example, how are two languages acquired, and the appropriate integrated-yet-differentiated representations established? To what extent are there critical periods or age of acquisition effects in the acquisition of an L2? To what extent are there transfer effects between first and second languages? How is an L2 best acquired – by initial association to an existing L1 or by a strategy that encourages direct contact with semantics (such as picture naming)? How is each language maintained, in terms of on-going patterns of relative dominance and/or proficiency, and to what extent are effects modality specific (i.e., differential across spoken and written language, comprehension and production)? How is each language lost, in terms of aphasia after brain damage in bilinguals, or in terms of the natural attrition of a disused language?

All of these questions demand a dynamic aspect of models of the bilingual system. But the point is not just about the features of bilingualism that remain to be explained. It also concerns the way one goes about constructing a model. For example, at times as I have watched the development of the BIA and now BIA+ framework, I have had a sense that the model is chasing the empirical data rather than predicting it. The authors have rarely put forward strong opinions about which features should be in the model and which not: whether there should be top-down connections or not; whether these connections should be inhibitory or not; whether interlingual homographs should have one or two representations; whether orthographic neighbourhoods should be facilitatory or inhibitory; whether phonological similarity should be inhibitory or facilitatory; and so on. There is little in the way of theory-driven expectation. Rather, data are presented and the model's assumptions altered accordingly. Indeed, sometimes the model's assumptions are altered between experiments, as in the case of the cross-language top-down inhibition in the BIA invoked to simulate one set of results but not another.

Now, this isn't necessarily a bad way of model building. Dijkstra, van Heuven and colleagues have made a great deal of progress building on existing static models from the monolingual literature. However, a complementary approach might seek to generate expectations based on what *function the relevant cognitive system has developed for*. What tasks is the system required to achieve in its normal usage? What representations is it therefore likely to learn to optimally succeed in these tasks (based on reasonable assumptions of the learning mechanisms available)?

These are not idle questions. A developmental perspective suggests theoretical possibilities for the structure of the *adult system* as well. Two examples will suffice, both theoretical issues that Dijkstra and van Heuven spend some time considering.

What is the role of language membership? The BIA/ BIA+ framework proposes that word nodes are present in an integrated network where each node has a link to a language node, which marks its membership. Part of the debate concerns the subsequent role of the activity that arrives at each language node. Inspection of the BIA implementation reveals that the connections between word nodes and language nodes are of identical value. Language membership is postulated to be a uniform construct. But is this really a sensible outcome of a plausible learning system? Let us say that the strength of the connection between a word node and its language node is established by Hebbian learning, where the language context establishes the activity on the language node at the same time as

the word node is activated, and their correlated activity strengthens the connection. Will every word node be associated with an equal degree of language context/language node activation? Will every word node be activated equally frequently? It seems unlikely.

In contrast, if we take a learning perspective, the expectation is that the system will evolve the processing structures sufficient to achieve the task at hand. In this case, the aim of the system is to provide the correct output (meaning) given the input (word form). Perhaps language context information will only be used in situations where it is necessary, for example, to disambiguate inputs that are similar across languages but require different outputs. Under this view, language membership may not be an all or nothing tag, but a continuum that depends on the demands of the task. Words with phonological or orthographic features unique to a language may not need to be tagged, whereas words with common features may lean more heavily on language membership information to disambiguate them.

(One might argue that all of a bilingual's words must have a language tag in order that the individual can answer the question "To which language does this word belong?" But such information could simply constitute encyclopaedic knowledge, part of the word's meaning. The extent to which language membership needs to assume an active *processing status* in the bilingual lexicon is a separate question.)

The second example follows on from this idea of the *constraints of learning*. Dijkstra and van Heuven debate whether interlingual homographs might have one or two representations in the bilingual lexicon, and whether cognates might have a special status. In constructing a static model, such distinctions must indeed take the form of *a priori* decisions. From a learning perspective, however, the system must find its own answer to this question, depending on the task that is being asked of it. If word form is to be identified, the representations need be no different for homographs/homophones, and increased frequency of exposure to the form across languages will likely improve response times and accuracy. But if the system is required to map from word form to semantics, interlingual homographs must have quite different representations, and likely must exploit language context information to drive this. Moreover, the degree to which each input form is associated with each meaning in the two languages is likely to have a material effect on the efficiency with which the system can generate each meaning; that is, the low frequency meaning of an interlingual homograph will be accessed more slowly (see Thomas, 1997 for a computational learning system exhibiting these characteristics).

In the same way that Dijkstra and van Heuven have built on existing static models of monolingual word recognition to account for phenomena within the bilingual domain, there now exists a host of monolingual *developmental* models addressing diverse phenomena. In the future, it will be possible to speak to many of the salient issues in bilingualism identified above, including issues of language acquisition, language dominance, language control and language loss (Thomas and van Heuven, in press).

The target article represents impressive progress in the development of theories of bilingual word recognition. In the future, one hopes that the models themselves will be permitted to develop too.

References

Thomas, M. S. C. (1997). Distributed representations and the bilingual lexicon: One store or two? In J. Bullinaria, D. Glasspool & G. Houghton (eds.), *Proceedings of the 4th Annual Neural Computation and Psychology Workshop*, pp. 240–253. Berlin: Springer.

Thomas, M. S. C. & van Heuven, W. J. B. (in press). Computational models of bilingual comprehension. In A. M. B. De Groot & J. F. Kroll (eds.), *The handbook of bilingualism: Psycholinguistic approaches*. Oxford: Oxford University Press.

AUTHORS' RESPONSE
Modeling bilingual word recognition: Past, present and future

TON DIJKSTRA
WALTER J. B. VAN HEUVEN
Nijmegen Institute for Cognition and Information (NICI),
Montessorilaan 3, P.O. Box 9104, 6500 HE Nijmegen,
The Netherlands.
E-mail: *dijkstra@nici.kun.nl, heuven@nici.kun.nl*

The detailed comments of our colleagues on our keynote article suggest that they all value the integration of various sorts of empirical data into one formalized model of bilingual word recognition. Taking this position, our peers propose further specifications or adaptations of the proposed BIA+ model with respect to its representations, processes or control structure; and they indicate what, in their view, should be the next step in model development or empirical testing. We are grateful for their contributions, which we will carefully think through in the coming time and use as a basis for future implementations.

In their commentaries, our colleagues of course do not only provide support for our model, but they also express different opinions about our modeling enterprise. We will discuss their major points of criticism in three sections, related to the following issues. First, commentators differ in their views of how one should model bilingual word recognition processes. Although they appear to agree that a formalized connectionist model of bilingual word recognition is useful, several believe that a distributed connectionist approach may be more correct or fruitful in the long run than the localist approach of the BIA+ model. Second, some authors provide interesting critical comments on certain model aspects in relation to available, sometimes recent, empirical data. We will respond to these comments, using them also to clarify aspects of the BIA+ model that have led to misunderstandings. Finally, commentators propose directions for future research and we will provide some first ideas where to go.

How to build a model of bilingual word recognition

As Li points out in his comments on the keynote article, there are so many variables that affect bilingual word recognition that one hardly knows where to begin to develop a formalized model. However, there are so many advantages to formalized (and implemented) models that, in spite of all difficulties, it is worth our while to make the attempt at capturing at least the most important aspects of word recognition in a formal model (see Dijkstra and De Smedt, 1996 for a discussion of the advantages of computational models). The BIA+ model, building on the BIA model, makes a step towards a more complete formalization including the major determining factors in word recognition although its computational implementation is still far from complete.

The basic problem of formalizing behavioral data from a particular subdomain of reality is, as Jerry Fodor (1987, Chapter 3) has expressed it, "how to carve nature at its joints". In other words, the question is how the subdomain should be simplified to ensure that the most important determinants of behavior are incorporated and interact correctly. The simulations of a successful model will mimic the behavior that is observed in the subdomain, and an examination of its internal architecture and processes will help us to better understand the behavior in terms of a complex interaction between multiple underlying causal mechanisms. While some researchers focus on important variables in adult lexical processing as a starting point for modeling, others consider the learning process and model dynamics as more important. Far from being unique to bilingual word recognition, this important issue has permeated cognitive science at large. Quite clearly, as Thomas points out, how one starts may affect where one ends, and focussing on one modeling aspect may lead to a relative neglect of another. Nevertheless, it seems useful to begin with a more simple model and then work towards a more complex one, and to build more complex models on the basis of earlier ones (nested modeling, see Jacobs and Grainger, 1994).

Each modeler must decide which aspects of the domain of interest should be taken as having to do with *structure* (relatively unchanging aspects of reality) and which with *process* (quickly changing aspects). This does not only hold for Interactive Activation (IA or localist connectionist) models, but also for Parallel Distributed Processing (PDP or distributed connectionist) models. For instance, PDP models are also restricted to learning in a particular subdomain (even though all aspects of reality are present simultaneously), and they presume the availability of representations such as letters or syllabic units (Dijkstra, in press). More generally, both IA and PDP models appear to have their strengths and weaknesses (see Grainger and Jacobs, 1998; Page, 2000 for the advantages of IA models).

Model development is stimulated by empirical research on the one hand, but it can drive empirical research on the other. Thus, as Li points out, modeling and experimentation should go hand in hand. However, in some areas, like bilingual sentence processing, so few data are available as yet that it is not obvious what it is exactly that we wish to model.

Are the BIA and BIA+ models only "chasing empirical data" (accounting for data only after they are collected), as Thomas puts it? We do not believe so. Admittedly, the BIA+ model has been adapted to account for recent empirical data, but it is based on explicit assumptions that generate new and testable predictions – as is seen, for instance, in the comments by Brysbaert, Van Wijnendaele and Duyck on our paper. We do not agree that we "have rarely put forward strong opinions about which features

should be in the model or not''. This becomes obvious if we examine Thomas's examples. In line with their IA modeling framework, the architecture of the BIA and BIA+ models from the very beginning implies the presence of interactions within the language processing system, e.g., between word representation levels. This interactivity can lead to top-down inhibition effects, the presence of which can be demonstrated empirically. (In other words, we have more of a problem to explain the *absence* of top-down inhibition from language nodes to words than its presence.) In addition, the BIA+ model explicitly states that non-linguistic context does *not* lead to top-down inhibition effects on the word identification system (note that hardly any connectionist models have concerned themselves with the relation to the task/decision system). This prediction stands in direct contrast to that of other views of language processing, such as the language mode hypothesis, the BIA model and the Inhibitory Control model.

Furthermore, the BIA/BIA+ models *do* predict the direction of orthographic neighborhood effects, but according to them this direction is dependent on the organizational characteristics of the lexicon(s) in question (see Van Heuven, Dijkstra and Grainger, 1998). As a consequence, the precise predictions cannot be made without taking the lexicon into account; they must be obtained by conducting model simulations with the lexicon(s). Such simulations can be considered as expressing the specific, context-sensitive predictions of the model. A similar point can be made for phonological similarity effects: depending on the orthographic and phonological characteristics of the lexicon involved, they will be facilitatory if particular (e.g., sub-lexical and lexical) units facilitate each other or if phonological activity can speed responses, and they will be inhibitory if (e.g., lexical) units compete for selection or response (see Jacobs, Rey, Ziegler and Grainger, 1998, for monolingual simulations with MROM-p). Here, again, simulations are needed to yield specific predictions, although knowledge of the interactive principles of IA models may help to formulate general expectations.

With respect to interlingual homographs, BIA+ explicitly assumes two orthographic representations. There appears to be some misunderstanding in how this choice was made. When we were developing the BIA model in collaboration with our colleague Jonathan Grainger, he rightly pointed out that the intuitively appealing implementation choice within the IA tradition was to have only one orthographic representation for interlingual homographs. (In some simulations we used two, but for practical reasons only, because BIA did not include phonological or semantic representations that might allow differentiation.) However, recent empirical data indicate that this option cannot be correct (see section 4.2 in the keynote article). Thus, "the only justification for this particular hand-coding" (Jacquet and French) is not that we as authors "know that 'ROOM' is an interlexical homograph", but it is the empirical evidence that we cannot ignore. Now, of course, the burden is on us to show that simulations with two representations will allow us to mimic empirical data involving interlingual homographs.

The BIA/BIA+ models make a distinction between lexical and language representations. In contrast to what Jacquet and French argue, words are *not* separated into different languages at the lexical level, but only at the higher level of the language nodes. All words are represented in the same way in a pool of items that is not itself differentiated by language. All words affect each other in the same way (through lateral inhibition), and all are connected to the same letter pools. Thus, BIA and BIA+ do not assume separate lexicons at the word level, but one integrated or unified bilingual lexicon. The misunderstanding may have arisen because of the figure of the model, in which English and Dutch words are labeled as such for the readers' sake.

The studies referred to by Colomé (2001) and Costa and Caramazza (e.g., 1999) on language production are compatible with our position. Indeed, these authors provide evidence supporting parallel lexical activation of word candidates in different languages and their proposed selection mechanism for word production has some commonalities with ours for comprehension. We do not agree with Jacquet and French that people know they are writing in English because an ensemble of English representations is "currently active in their brains, period". Word production research shows that lexical candidates from both languages of the bilingual initially become active and language membership can serve as a selection cue for filtering out words from the non-target language (Kroll, Dijkstra, Janssen and Schriefers, 1999; Kroll and Dijkstra, 2002).

In the BIA/BIA+ models, word-specific language membership information becomes available only via the word level. Language information is less important in BIA+ because it is purely representational and provides a weak source of constraint on word recognition only. Nevertheless, the language labels are assumed to become available to the reader at only a short delay relative to word identification. In contrast to Thomas, we believe that the knowledge that a word belongs to a particular language is not merely encyclopaedic but useful to and used by the bilingual. For instance, we assume that in language decision participants use the language representation to make their decision which button to press, and in order to retrieve the item's language membership, they have to identify the item.

In contrast to Jacquet and French' opinion, the BIA+ and BIA models can account quite well for between-language neighborhood effects even without resorting to the language nodes (Dijkstra and Van Heuven, 1998). The mechanisms of lateral inhibition between word nodes, adaptations of the resting level activation of words depending on subjective frequency, and possible adaptations of decision criteria, appear to be sufficient to account for the data patterns observed by French and Ohnesorge (1996) and ourselves (Van Heuven et al., 1998).

As Thomas points out, the BIA and BIA+ connections between word nodes and language nodes are of identical value, making language membership what he calls "a uniform construct". Indeed, in an IA approach, the connection weights between representations of different levels are usually (but not necessarily) hand-set at one normalized value (for instance, depending on the total number of connections between levels). This practice follows the prin-

ciple of *parsimony*, and it considerably reduces the number of model parameters. The connection strengths for different words are usually equated at a mean value under the assumption that ignoring their differences will not systematically affect average responding. In the present case, this leads to an interesting prediction. If the weights between representation levels should be set at different values, as Thomas suggests, latencies for low frequency words should increase more than for high frequency words in tasks where the response is read out from the language node level relative to tasks where it is read out from the lexical level. This is not predicted in case all weights between levels are set equally. However, it may be that this hypothesis can only be tested empirically if additional assumptions are made.

Empirical aspects of the BIA+ model

Following the IA model of monolingual word recognition, the BIA model made the structural choice of incorporating only letter and word representations, and no intermediate representations such as syllables. For BIA+ we are still experimenting with sublexical units larger than letters or phonemes, e.g., Onset–Nucleus–Coda (ONC) units (Van Heuven, 2000). Obviously, in order to do any modeling at all, a choice has to be made, but any choice may be at the expense of model flexibility and the ability to account for particular data sets (Dijkstra, in press). In line with Green's argumentation, certain grain-size effects may be impossible to simulate without assuming particular intermediate levels (see Martensen, Maris and Dijkstra, 2000, for evidence that the grain size of word recognition units may be different across languages).

An interesting question is which grain-size effects will turn out to be dependent on lexical characteristics (orthography, phonology) and possible to simulate, and which will be out of the model's range (e.g., certain strategic effects). For example, we found that certain word body effects (Ziegler and Perry, 1998) emerged in a simulation with a monolingual model incorporating ONC units, even though word body representations were not explicitly included. The effect appeared to be a consequence of the organization of the orthographic and phonological lexicons used by the model. However, given the simplifications that will have to be made, the present architecture of the BIA+ model will almost certainly turn out to be too simple to capture certain subtle cross-linguistic or cross-modal effects. (Thus, it may be necessary to develop different models for different purposes.)

Grain size differences also play a role with respect to the involvement of sublexical and lexical phonology in bilingual word recognition. According to the "temporal delay" hypothesis, L2 phonological and semantic codes are delayed in activation relative to L1 codes, depending on the proficiency of the bilingual. Evidence supporting this view was found in the study by Lemhöfer and Dijkstra (submitted). We wish to make two remarks with respect to the temporal delay hypothesis. First, the hypothesis is *lexical* in nature, having to do with the activation of lexical rather than sublexical representations. With respect to *sublexical*

representations, the BIA+ model assumes that words from different languages having the same script share their letters and some of their phonemes (cf. Roelofs, in press). As a consequence, sublexical cross-linguistic priming effects might be independent of word frequency, in line with the recent study by Van Wijnendaele and Brysbaert (2002).

Second, even though the lexical-phonological and semantic L2 codes are generally delayed relative to L1 codes, the effect depends on the subjective frequency of the item in question. For instance, for interlingual homographs with low frequency Dutch readings and high frequency English readings, English phonology or semantics could still be retrieved earlier in time. Thus, the delay will not be constant for all types of words, as Brysbaert et al. assume. Notwithstanding these remarks, we take to heart Brysbaert et al.'s *caveat* that interactions between cross-linguistic codes may be more dynamic than implied by the term "temporal delay assumption". It will be interesting to see the limits to which the lexical hypothesis can be pushed.

Several colleagues pointed out that our empirical research and modeling has so far paid relatively little attention to semantic/conceptual issues (Kroll and Dussias, in press; Van Hell; Jacquet and French; and Brysbaert et al.). That is true. We have been doing preliminary simulations involving a semantic association network (Van Heuven, 2000), but this work has only just started. On the empirical side, we have recently obtained cross-linguistic effects of morphological family size on interlingual homographs in collaborative work with Baayen (MPI) and Schreuder's (IWTS) research group. The morphological family of a word consists of those morphologically related words that contain the word as a member (e.g., "appeltaart" is in the family of "appel", see De Jong, 2002). Morphological family size effects appear to be at least partly semantic in nature. Our analyses indicate that Dutch-English interlingual homographs are affected in an opposite way by the size of their morphological families in both languages. Thus, evidence for yet another independent variable shows that bilingual word recognition is thoroughly language non-selective.

Referring to studies mentioned by Kroll and Dussias (in press), Van Hell rightly points out that the explanation of (in particular) syntactic effects of sentence context on the recognition of target words in bilinguals may require incorporating a whole new apparatus into the BIA+ model. We agree with her that "language-specific properties may challenge the universality and language-independence of processing routines". However, we note that the research rationale so far has been to test the non-selective access view for processing domains where non-selectivity is possible at all. For language-specific phenomena (at sublexical, lexical and syntactic levels), language-specific results are bound to be found. However, even for language non-specific phenomena the issue may be quite complex, as is shown by a pilot experiment conducted in our lab (Caelen, 1998). When Dutch participants processed Dutch (L1) sentences with temporarily ambiguous subject- and object-relative clauses, they appeared to use mainly syntactic sources of information to resolve the ambiguity. The same held for German participants performing in their L1,

German. In contrast, the performance of Dutch-German bilinguals with respect to German (L2) sentences was affected by both syntactic and semantic information (a finding that has also been reported for monolinguals under more demanding circumstances). Possibly, the larger processing load for L2 sentence processing in bilinguals led to performance differences relative to L1 processing that were both quantitative (latencies) and qualitative (effects of syntax and semantics).

Roelofs considers which role a non-linguistic factor such as attention can play in experiments if it does not directly modulate activation in particular processing pathways. As an alternative, he suggests that the attention can be focussed on those production rules in the task schema that include the readers' goal among their conditions. As far as we can see, goal-referenced control as he envisions it can readily be implemented in the BIA+ model, also because the results of bilingual Stroop tasks appear to be directly compatible with BIA+.

Green points out that readers might become sensitive to subtle sublexical language-specific sources of information in the input word, in order to speed their responses, for instance, in a language decision task. Studies on cue validity indeed suggest that (even unaware) participants in experiments may sometimes exploit probabilistic relationships between stimulus and response. However, whether this holds for sublexical information correlated with language membership remains to be empirically (re)-confirmed (cf. Grainger and Beauvillain, 1987). Furthermore, because this proposed mechanism operates only bottom-up, it could in principle be implemented in the BIA+ framework.

Green further describes a mechanism that accounts for carry-over effects between trials that we believe to be compatible with the BIA+ model. In this context, he refers to the notion of "full activation", a term we used in our paper. However, we apparently were not clear enough in our explanation of the term. We do not believe that "even when a speaker has not used a language for years, the lexical representations of that language compete, just as strongly as they ever did, with the current language to account for the output". The resting activation levels of representational units are assumed to decrease if a language is not used for a while (Dijkstra and Van Hell, in press). As a consequence, such representations will experience more lateral inhibition exerted by their lexical competitors during recognition. The accompanying changes in the pattern of activation of the various competitors will affect the moment in time the response is made.

Finally, we note that the explicit distinction between the word identification system and the task/decision system facilitates our understanding of how performance differs across experimental situations. This approach clarifies which elements of the task and recognition process are affected differently in different tasks. In our view, this provides a strong *advantage for the modeler* relative to approaches that merely inform us of the priority of contributing factors in different situations or that provide different matrices of learned connection weights (cf. Thomas).

Future work

In their comments, our colleagues point to important empirical domains that remained unmentioned, such as second-language acquisition and learning (Li; Thomas; Jacquet and French), interindividual differences (Green; Van Hell), and the relationship between functional modeling and the brain (Green). As any researcher will understand, we have concentrated our research and modeling efforts on those areas where our expertise is greatest. It is, therefore, at the risk of sounding naïve that we offer some thoughts about potential model development in these areas.

With respect to acquisition issues, we would like to point out that the BIA/BIA+ framework appears to be quite compatible with and complementary to recent versions of the Revised Hierarchical Model (RHM) by Kroll and her colleagues (e.g., Kroll and Sunderman, in press). This model incorporates an account of second-language acquisition not present in the BIA+ model, and several of the questions raised by Thomas could be answered in terms of the RHM (see Hulstijn, 2002 for further considerations on second-language acquisition). As one example of the compatibility of the BIA+ model and the RHM, we may consider that the RHM hypothesized that in early L2 learners, translation equivalents were active at the lexical-association level (e.g. stoel [D] – chair [E], while the BIA+ focussed on the cross-linguistic effect of form relatives (e.g., stoel [D] – steel [E]). Recent evidence by Sunderman and Kroll suggests that both effects are present. Form relatives are active for learners at both early and late stages of second-language acquisition, just as BIA+ predicts. However, translation associates may be active only for the learners at the earliest stages of L2 acquisition. Changes in form activation due to increases in L2 proficiency may be accounted for in the BIA+ model by adjusting the resting level activations (reflecting subjective frequency), but this will not account for the presence or absence of translation activity. However, BIA+ does not disagree with the RHM here, simply because additional assumptions about lexical association and semantics are needed in order to account for the effects observed for translation equivalents.

With respect to more general aspects of learning, we point out that, even though distributed (PDP) connectionist models are well-known for their learning capabilities, localist connectionist models such as BIA can be made to learn as well. A good example is the work by Grossberg (1987), who introduced a learning mechanism in his Adaptive Resonance Theory (ART); also see Happel and Murre, 1994; Burton, 1998).

Even without modeling second-language learning processes or increases in L2 proficiency directly, one can simulate bilingual word recognition performance for bilinguals at different proficiency levels. For instance, to simulate the lexical decision data for two high-school groups, we will adapt the L2 word frequencies in the BIA+ model separately for each proficiency group (Schulpen, Dijkstra and Schriefers, in preparation).

Most simulations with the BIA model (including those still valid for BIA+) have tried to mimic patterns of RTs averaged across participants or item categories. However,

in principle some or even all parameters of the model can be fit individually for each participant or item. Thus, some interindividual differences can be simulated by the BIA+ model. However, it will require extensive experimentation and complex parameter setting procedures to fit the data. Parameters that could be varied are resting level activation (reflecting subjective frequency), bottom-up activation and inhibition between letter and word levels, and decision criteria (e.g., amount of activation required for "yes" responses).

Finally, Green brings up the complex and general issue of the relationship between neuroimaging data and functional models, of which BIA+ is only one example. In our view, the fast development of cognitive neuroscience is exciting and it is our conviction that any complete model of bilingual language processing should ultimately take into account reliable data from multiple sources, including brain activity data and behavioral data.

However, at present, the relationship between functionally distinct (soft) modules and localization in the brain is not well understood. An important insight, we think, is that functional models are based on behavioral data and are not models of brain functioning. There appears to be a many-to-many mapping between components in our functional models and areas in the brain. On the one hand, one function can be associated with activity in different parts of the brain. This already holds for language processing as a whole, which (although functionally differentiated) takes place in both hemispheres (Federmeider and Kutas, 1999). At a more specific level, Kaan and Swaab (2002) argue that syntactic processing does not recruit one specific area in the brain but that a whole network of areas is involved. On the other hand, one brain area may be engaged in various functions, as Green points out in his commentary. This observation makes it particularly hard to allocate a psychological interpretation to certain electrophysiological and neuroimaging data. For instance, with respect to ERP-measurements the P600 component was assumed to be a marker in the EEG associated with syntactic violations, but recently this interpretation has been challenged (Kolk, Chwilla, Van Herten and Oor, submitted).

In sum, systematic patterns of brain activity need not directly correspond to the conceptual entities in our functional models. Observations like these suggest that linguistic functions shown to be distributed in the brain need not be distributed in our functional models for behavior. However, a demonstration that different brain areas are systematically engaged in different subfunctions, might be followed by the collection of empirical evidence for a related distinction in behavioral data and then by changes in the model.

Given the current state of affairs, we do not find it surprising that different studies may collect apparently paradoxical results at least for some time to come. An example pointed out by Roelofs and Van Hell is the recent ERP/fMRI study by Rodrigues-Fornells, Rotte, Heinze, Nösselt and Münte (2002), which is interpreted by the authors as evidence that "words from the non-target language are rejected at an early stage before semantic analysis in bilinguals". Roelofs argues that the findings of this study do not imply controlled language-specific access. He points out that reduced effects for words from the non-target language may lead to an absence of effects on brain potentials, and that bilinguals process words from two languages in contrast to the monolinguals.

We would like to add a few points of our own. First, the conclusion was partially based on the ERPs for *non-target* words and pseudowords, items that did not require explicit responses. Thus, it is questionable if the conclusion may be generalized to the target words. In fact, in a recent ERP study of our own (Kerkhofs, Dijkstra, Chwilla and De Bruijn, in preparation), the N400s for interlingual homographs *were* sensitive to the relative frequency of their readings in both the target and non-target language. Second, the authors' conclusion leaves open the possibility of parallel activation of *orthographic* lexical candidates from different languages. Third, the response hand for the go-stimuli in the study depended on the first letter (vowel or consonant) of the stimulus, which complicated the participants' task, possibly leading to strategy changes and to extra activation of sublexical phonological pathways. In sum, it would be premature to return to a version of the language-selective access view that ignores the many RT studies supporting parallel activation of word candidates at the semantic/conceptual level.

In spite of these and other challenges, we must continue to relate our functional models to brain imaging data. For instance, the distinction between the word identification system and the task/decision system proposed by Green's Inhibitory Control model and the BIA+ model could be reflected in the brain by activity in the temporal lobe (e.g., the posterior mid-fusiform gyrus or word form area referred to by Green) and activity in the anterior cingulate cortex (cf. Roelofs and Hagoort, submitted). We have planned to test this view in an event-related fMRI variant of the experiments by Dijkstra, Van Jaarsveld and Ten Brinke (1998). The relative brain activity in the three experiments might support the view that the linguistic processing aspects remain similar across experiments, while the task demands and control aspects change.

References

Burton, A. M. (1998). A model of human face recognition. In J. Grainger & A. M. Jacobs (eds.), *Localist connectionist approaches to human cognition*, pp. 75–100. Mahwah, NJ: Lawrence Erlbaum Associates.

Caelen, M. (1998). *Extending the study on the processing of relative clauses to bilingualism*. Master's thesis, University of Nijmegen.

Colomé, A. (2001). Lexical activation in bilinguals' speech production: language-specific or language-independent? *Journal of Memory and Language*, 45, 721–736.

Costa, A. & Caramazza, A. (1999). Is lexical selection in bilinguals language specific? Further evidence from Spanish-English bilinguals. *Bilingualism: Language and Cognition*, 2, 231–244.

De Jong, N. H. (2002). Morphological families in the mental lexicon. Ph.D. thesis, University of Nijmegen.

Dijkstra, A. (Ton) (in press). Word recognition and lexical access II: Connectionist approaches. In D. A. Cruse, F. Hundsnurscher, M. Job & P. R. Lutzeier (eds.), *Lexikologie – Lexicology*. Berlin: Walter de Gruyter. (Article 265.)

Dijkstra, A. (Ton) & De Smedt, K. (1996). Computer models in

psycholinguistics: An introduction. In A. (Ton) Dijkstra & K. De Smedt (eds.), *Computational psycholinguistics: AI and connectionist models of human language processing*, pp. 3–23. London: Taylor & Francis.

Dijkstra, A. (Ton) & Van Hell, J. G. (in press). Testing the Language Mode hypothesis using trilinguals. *International Journal of Bilingual Education and Bilingualism*.

Dijkstra, A. (Ton) & Van Heuven, W. J. B. (1998). The BIA model and bilingual word recognition. In J. Grainger & A. M. Jacobs (eds.), *Localist connectionist approaches to human cognition*, pp. 189–225. Mahwah, NJ: Lawrence Erlbaum Associates.

Dijkstra, A. (Ton), Van Jaarsveld, H. & Ten Brinke, S. (1998b). Interlingual homograph recognition: Effects of task demands and language intermixing. *Bilingualism: Language and Cognition*, 1, 51–66.

Federmeier, K. D. & Kutas, M. (1999). Right words and left words: Electrophysiological evidence for hemispheric differences in meaning processing. *Cognitive Brain Research*, 8, 373–392.

Fodor, J. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press/A Bradford Books.

French, R. M., & Ohnesorge, C. (1996). Using interlexical non-words to support an interactive-activation model of bilingual memory. *Proceedings of the Eighteenth Annual Cognitive Science Society Conference*, pp. 318–323. Mahwah, NJ: Lawrence Erlbaum Associates.

Grainger, J. & Beauvillain, C. (1987). Language blocking and lexical access in bilinguals. *Quarterly Journal of Experimental Psychology*, 39A, 295–319.

Grainger, J. & Jacobs, A. M. (1998). On localist connectionism and psychological science. In J. Grainger & A. M. Jacobs (eds.), *Localist connectionist approaches to human cognition*, pp. 1–38. Mahwah, NJ: Lawrence Erlbaum Associates.

Grossberg, S. (1987). Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science*, 11, 23–63.

Happel, B. L. M. & Murre, J. M. J. (1994). Design and evolution of modular neural networks. *Neural Networks*, 7, 985–1004.

Hulstijn, J. (2002). Towards a unified account of the representation, processing and acquisition of second language knowledge. *Second Language Research*, 18, 193–223.

Jacobs, A. M. & Grainger, J. (1994). Models of visual word recognition: Sampling the state of the art. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 1311–1334.

Jacobs, A. M., Rey, A., Ziegler, J. C. & Grainger, J. (1998). MROM-p: An interactive activation, multiple readout model of orthographic and phonological processes in visual word recognition. In J. Grainger & A. M. Jacobs (eds.), *Localist connectionist approaches to human cognition*, pp. 147–188. Mahwah, NJ: Lawrence Erlbaum Associates.

Kaan, E. & Swaab, T. Y. (2002). The brain circuitry of syntactic comprehension. *Trends in Cognitive Sciences*, 6, 350–356.

Kerkhofs, R., Dijkstra, A. (Ton), Chwilla, D. J. & De Bruijn, E. (in preparation). How relative word frequency affects RTs and ERPs to interlingual homographs in bilingual semantic priming. [working title]

Kolk, H. H. J., Chwilla, D. J., Van Herten, M. & Oor, P. J. W. (submitted). Structure and limited capacity in verbal working memory: A study with event related potentials.

Kroll, J. F. & Dijkstra, A. (Ton) (2002). The bilingual lexicon. In R. Kaplan (ed.), *Handbook of applied linguistics*, pp. 301–321. Oxford: Oxford University Press.

Kroll, J. F., Dijkstra, A. (Ton), Janssen, N. & Schriefers, H. (1999). *Cross-language lexical activity during production: Evidence from cued picture naming*. Paper presented at the European Society for Cognitive Psychology, Ghent, Belgium, September.

Kroll, J. F. & Dussias, P. E. (in press). The comprehension of words and sentences in two languages. In T. Bhatia & W. Ritchie (eds.), *Handbook of bilingualism*. Cambridge MA: Blackwell Publishers.

Kroll, J. F. & Sunderman, G. (in press). Cognitive processes in second language acquisition. In C. Doughty & M. Long (eds.), *Handbook of second language acquisition*. Cambridge, MA: Blackwell Publishers.

Lemhöfer, K. & Dijkstra, A. (Ton) (submitted). Recognizing cognates and interlingual homographs: Effects of code similarity in generalized lexical decision.

Martensen, H., Maris, E. & Dijkstra, A. (Ton) (2000). When does inconsistency hurt? On the relation between phonological-consistency effects and the reliability of sublexical units. *Memory & Cognition*, 28, 648–656.

Page, M. (2000). Connectionist modeling in psychology: A localist manifesto. *Behavioral and Brain Sciences*, 23, 443–512.

Rodrigues-Fornells, A., Rotte, M., Heinze, H.-J., Nösselt, T. & Münte, T. (2002). Brain potential and functional MRI evidence for how to handle two languages with one brain. *Nature*, 415, 1026–1029.

Roelofs, A. (in press). Shared phonological encoding processes and representations of languages in bilingual speakers. *Language and Cognitive Processes*.

Roelofs, A. & Hagoort, P. (submitted). Control of language use: Cognitive modeling the hemodynamics of Stroop task performance.

Schulpen, B., Dijkstra, A. (Ton) & Schriefers, H. J. (in preparation). A cross-sectional study on the development of automatization and cognitive control in Dutch-English bilinguals: Part 1. [working title]

Van Heuven, W. J. B. (2000). Visual word recognition in monolingual and bilingual readers: Experiments and computational modeling. Ph.D. thesis, University of Nijmegen. (NICI Technical Report 20-01.)

Van Heuven, W. J. B., Dijkstra, A. (Ton) & Grainger, J. (1998). Orthographic neighborhood effects in bilingual word recognition. *Journal of Memory and Language*, 39, 458–483.

Van Wijnendaele, I. & Brysbaert, M. (2002). Visual word recognition in bilinguals: Phonological priming from the second to the first language. *Journal of Experimental Psychology: Human Perception and Performance*, 3, 619–627.

Ziegler, J. C. & Perry, C. (1998). No more problems in Coltheart's neighborhood: Resolving neighborhood conflicts in the lexical decision task. *Cognition*, 68, B53–B62.