Cole, P., Beauvillain, C., Pavard, B., & Segui, J. (1986). Organisation morphologique et acces au lexique. *L'Annee Psychologique, 86,* 349-365.
Cole, P., Beauvillain, C., & Segui, J. (1989). On the representation and processing of prefixed and suffixed derived words: A differential frequency effect. *Journal of Memory and Language, 28,* 1-13.
Cutler, A., Hawkins, J. A., & Gilligan, G. (1985). The suffixing preference: A processing explanation. *Linguistics, 23,* 723-758.
Forster, K. I. (1987). Form-priming with masked primes: The best match hypothesis. In M. Coltheart (Ed.), *Attention and performance XII.* Hillsdale, NJ: Erlbaum.
Forster, K. I., & Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. *Journal of Experimental Psychology: Learning, Memory and Cognition, 10,* 680-690.
Forster, K. I., Davis, C., Schoknecht, C., & Carter, R. (1987). Masked priming with graphemically related forms: Repetition or partial activation? *Quarterly Journal of Experimental Psychology, 39A,* 211-251.
Fowler, C. A., Napps, S. E., & Feldman, L. (1985). Relations among regular and irregular morphologically related words in the lexicon as revealed by repetition priming, *Memory & Cognition, 13,* 241-255.
Grainger, J., Cole, P., & Segui, J. (1991). Masked Morphological Priming in Visual Word Recognition. *Memory and Language, 30,* 370-384.
Hawkins, J. A., & Cutler, A. (1988). Psycholinguistic factors in morphological asymmetry. In J. A. Hawkins (Ed.), *Explaining language universals.* Oxford, Basil Blackwell.
Henderson, L., Wallis, J., & Knight, D. (1984). Morphemic structure and lexical access. In H. Bouma & D. G. Bouwhuis (Eds), *Attention and performance X.* London, Erlbaum.
Henderson, L. (1985). Toward a psychology of morphemes. In A. W. Ellis (Ed.), *Progress in the psychology of language* (Vol 1). London, Erlbaum.
Lukatela, G., Gligorijević, B., Kostić, A., & Turvey, M.T. (1980). Representation of inflected nouns in the internal lexicon. *Memory & Cognition, 8,* 415-423.
Manelis, L., & Tharp, D. A. (1977). The processing of affixed words. *Memory & Cognition, 5,* 690-695.
Murrell, G. A., & Morton, J. (1974). Word recognition and morphemic structure. *Journal of Experimental Psychology, 102,* 963-968.
Napps, S. E. (1989). Morphemic relationships in the lexicon: Are they distinct from semantic and formal relationships? *Memory & Cognition, 17,* 729-739.
Schriefers, H., Zwitserlood, P., & Roeloffs, A. (1991). The identification of morphologically complex words: Continuous processing or decomposition? *Journal of Memory and Language, 30,* 26-47.
Napps, S. E., Fowler, C. A. (1987). Formal relationships among words and the organization of the mental lexicon. *Journal of Psycholinguistic Research, 16,* 257-272.
Segui, J., & Grainger, J. (1990). Priming word recognition with orthographic neighbors: Effects of relative prime-target frequency. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 65-76.
Segui, J., & Zubizarreta, M-L. (1985). Mental representation of morphologically complex words and lexical access. *Linguistics, 23,* 759-774.
Taft, M. (1985). The decoding of words in lexical access: A review of the morphographic approach. In D. Besner, T. G. Waller, & G. E. Mackinnon (Eds.), *Reading research: Advances in theory and practice* (Vol 5). London: Academic Press.
Tyler, L. K., Marslen-Wilson, W., Rentoul, J., & Henney, P. (1988). Continuous or discontinuous access in spoken word recognition: The role of derivational prefixes. *Journal of Memory and Language, 27,* 368-381.

Orthography, Phonology, Morphology, and Meaning
R. Frost and L. Katz (Editors)

CHAPTER 20

Bilingual Lexical Representation: A Closer Look at Conceptual Representations

Annette M. B. de Groot
Department of Psychology
University of Amsterdam, The Netherlands

Although usually extremely simple, the figures that illustrate bilingual memory organization in journal articles and book chapters often do a good job of accounting for the data. Yet, when studying them, one cannot help wondering every so often whether, rather than parsimoniously capturing its essence, these few strokes and dashes may do injustice to the complexity of reality. Take, as an example, Figure 1. It is based on my own work on between-language repetition and semantic priming (de Groot & Nas, 1991) and on word translation (de Groot, in press).

As is often done when depicting bilingual memory organization, two representational levels are distinguished. A whole word is represented in a single node at the lexical level; its meaning in a single node at the conceptual level. In other papers (e.g., Chen & Leung, 1989; Potter, So, Von Eckardt, & Feldman, 1984), instead of circles for individual words and their meanings, boxes are drawn to represent whole word and concept systems.

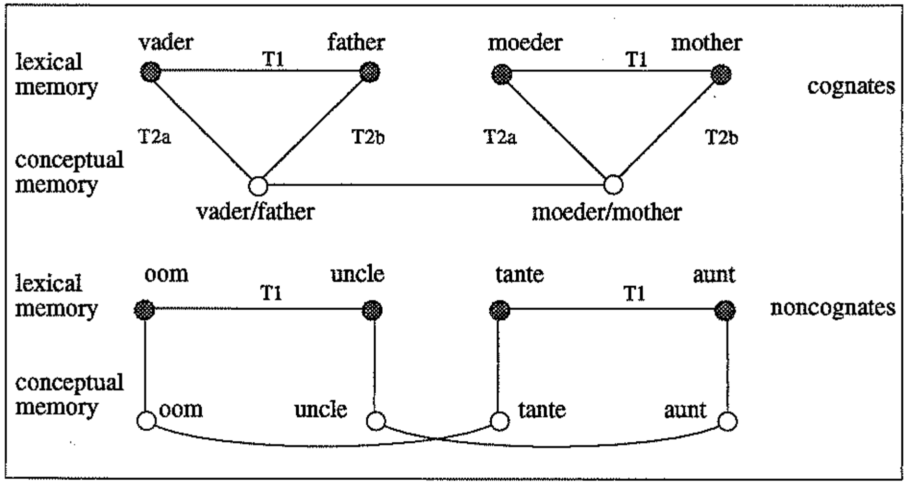

Figure 1. **Some representations in bilingual memory.**

What I primarily intend to do here is to zoom in on some of the circles of Figure 1 to see what can be discerned there. In so doing, the various parts of Figure 1 will be discussed, and standard accounts of a number of results from bilingual investigations will be reviewed. In addition, although they are not backed by new data, alternative explanations will be suggested.

Although their content was never explicated, the circles at the conceptual level were never deliberately intended to suggest indivisible entities. In Quillian's hierarchic network model of semantic memory (e.g., Collins & Quillian, 1969; Quillian, 1968), from which many of the views on monolingual and bilingual representation are derived, concepts were represented in nodes, relations between concepts in links between nodes, and the meaning of a concept by the pattern of relationships in which the concept node participates (see Rumelhart & Norman, 1985). So the meaning of *bird* would consist of 'is a subset of *animal,*' 'has as subset *canary*', 'has as subset *ostrich,*' 'has as parts *feathers,*' 'can fly,' etc. Or in the non-hierarchical successor of this model, the meaning of *red* would consist of its relation with *orange, yellow, green, fire, apples, roses,* etc. (Collins & Loftus, 1975). In the same vein, the concept nodes in Figure 1 can be seen as built up from a number of meaning elements. This is made explicit in Figure 2. The concept associated with the word *vader (father)* in Figure 1 is now spread out over six nodes, each of them representing one meaning element of the word *vader*. The number six is chosen arbitrarily. I will henceforth call these conceptual representations 'distributed' (see e.g., Hinton, McClelland, & Rumelhart, 1986). Instead of there being just one connection from the lexical node for *vader* to its conceptual node (Figure 1), the lexical node now has connections to each of the meaning elements of the conceptual representation. Upon presentation of the word *vader*, each of these elements receives excitatory activation via its connection with the lexical node.

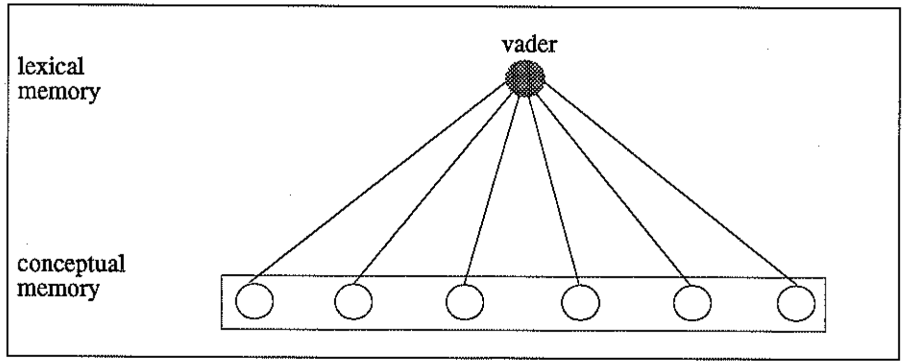

Figure 2. **A distributed conceptual representation in memory.**

In the next parts of this chapter I will gradually expand Figure 2 with the purpose of providing new accounts of some of the results obtained in a number of bilingual (and occasionally monolingual) processing tasks (the essence of some of these accounts was also suggested by Taylor and Taylor (1990). Only semantic memory tasks will be considered, that is, tasks that *could* be performed without consulting episodic knowledge (even though performance may well be influenced by such knowledge). The tasks being discussed are: word translation, primed lexical decision (that is, intra- and interlingual semantic priming and repetition priming, with lexical decision serving as the experimental task), word association, and semantic-relation assessment.

Word translation

In a number of studies the word translation task has been used as a means of obtaining information on the organization of knowledge in bilingual memory (e.g., Chen & Leung, 1989; de Groot, in press; Kroll & Curley, 1988; Potter et al., 1984). In its standard form the task simply involves presenting the bilingual subjects with words in one language, and asking them to produce their translation in a second language. In all but one of these studies (de Groot, in press) the task was used in conjunction with one or more other tasks, most often picture naming in a second language. A comparison of response times in word translation and picture naming in the second language was meant to solve the question of whether word translation takes place via a direct connection between the lexical representations of the translation equivalents (Route T1 in Figure 1), or indirectly, via an amodal conceptual representation shared by the two translation equivalents as well as by a picture of the referent of these words (Route T2 [T2a + T2b] in Figure 1; the node for the picture is not shown). If translation comes about by tracing T1, it is argued, translation should take less time than picture naming in the second language, because the route to the response would be shorter than in picture naming (in which access of the conceptual node cannot be circumvented). But if Route T2 is traced in translation, word translation and picture naming should take equally long. In the case of the latter outcome one might want to conclude, as Potter et al. (1984) did, that no direct connections exist between the representations of translation equivalents at the lexical representational level.

Fluent bilinguals turn out to be as fast in second-language picture naming as in word translation (Chen & Leung, 1989; Kroll & Curley, 1988; Potter et al., 1984), but less proficient bilinguals (or, more precisely, less proficient *adult* bilinguals; Chen & Leung, 1989) translate faster than they name pictures in their second language (Chen & Leung, 1989; Kroll & Curley, 1988). The data thus indicate that fluent bilinguals use Route T2, whereas less proficient adult bilinguals take Route T1. This suggests that T1-connections do exist, but are bypassed by fluent bilinguals during word translation. From a study comparing naming in a native (Dutch) and a second (English) language, on the one hand, with translation between the two languages on the other, Kroll and Stewart (1990) indeed drew the conclusion that T1-connections exist (see also de Groot & Nas, 1991). At the same time they qualified this conclusion: There are T1-connections in both directions, from the stronger to the weaker language and vice versa, but they differ in strength. The link from the weaker (here English) to the stronger language (here Dutch) is the stronger of the two.

The process involved in what was called 'tracing' translation routes above is presumably 'spreading activation': When a stimulus word is presented, it first contacts its representation in lexical memory. The activation that originates in this memory node spreads out along the paths of the memory network and activates the representations it encounters en route. Consequently, the activation levels of the encountered representations are temporarily increased, and the corresponding words are readily available as responses. Which of these activated representations is eventually selected for responding will depend on the task at hand (and on the extent to which each single one of them is activated). In word translation by fluent bilinguals, the representation selected for responding will typically be that of the stimulus word's translation in the target language.

How can word translation by fluent bilinguals be depicted in light of the representations illustrated in Figure 2? Figure 3a shows how this can be done. It repeats Figure 2, but now a lexical node for the English translation equivalent of *vader* is added. Additionally, links between this representation and each of the meaning elements constituting the conceptual representation are shown. The implicit assumption in this example, right or wrong, is that the Dutch word and its English translation have exactly the same meaning. Translation again involves the tracing of links (spreading activation) from the lexical representation of a word in one language to that of this word in the second language via conceptual memory, but now the links to be traced are those connecting the lexical nodes with the individual meaning elements of the conceptual representation.

The meanings of translations often do not fully overlap. This may more often be the case with some types of words than with others. For instance, it has been suggested that the meanings of abstract words differ more across languages than those of concrete words (Taylor, 1976). Such a state of affairs can be captured in a very straightforward way by representational structures of the kind depicted in Figure 3a. Figure 3b shows a situation wherein an abstract word in Dutch shares four meaning elements with the corresponding word in English. In addition to the common meaning elements, each word has two not shared by the corresponding word in the other language. Thus, the conceptual representations now overlap only partly.

It is reasonable to assume that the situations depicted in Figures 3a and 3b lead to differences in translation performance: The more conceptual elements shared by a pair of translation equivalents, the more activation will spread from the lexical node of a word to that of its translation, and the better performance—as assessed by response speed, number of errors, and number of 'omissions' (where no translation is given)—will be. Staying with our examples, other things being equal, the concrete word *vader* should be translated faster and/or more often and/or more often correctly than the abstract word *idee*.

I indeed observed these differential effects for concrete and abstract words in a recent study on word translation (de Groot, in press), although I argued there that another word characteristic than word concreteness, namely, 'context availability' (see e.g., Schwanenflugel, Harnishfeger, & Stowe, 1988) may well underlie them (but this does not undermine the point I am trying to make here; I could have chosen another word characteristic than word concreteness to illustrate it). In this same study I also obtained differential effects for cognates (words with orthographically and phonologically similar translations) and noncognates (dissimilar translations). Performance was consistently better for the former type of words.

In sum, it appears that a memory with distributed conceptual representations could quite naturally cope with differential translation performance on different types of words.

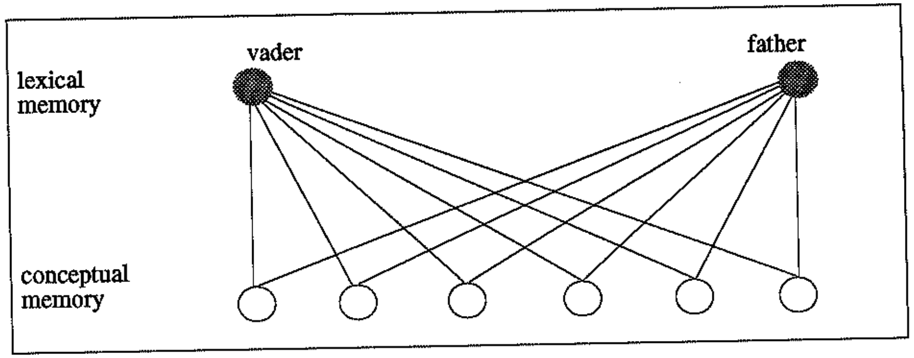

Figure 3a. **A distributed conceptual representation in memory. Translations have exactly the same meaning.**

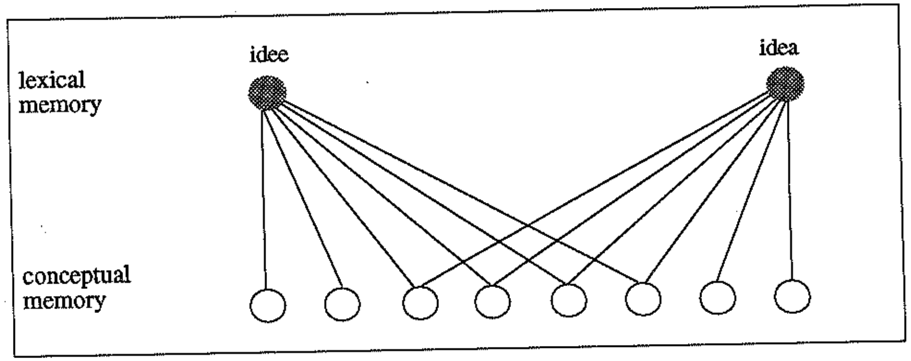

Figure 3b. **Distributed conceptual representations in memory. Translations differ in meaning.**

Note, however, that I am not suggesting that effects of *all* word characteristics on translation performance could be explained this way. For instance, word frequency effects (de Groot, in press) may be due to differences in the strength of connections (e.g., between the lexical and conceptual nodes in Figure 1, or between the lexical nodes and the nodes for the various meaning elements in Figures 3a and 3b) between words of high and low frequency, rather than to different numbers of shared meaning elements between the translations. The stronger the links, the more activation they transmit.

Errors in word translation

Words share aspects of meaning not only with their translations, but also with semantically related words in the same language. They may also share parts of their meaning with the translations of these semantically related words. Figures 4a and 4b serve to illustrate this point. At the risk of being too explicit I add that the amount of semantic overlap between the various words suggested in these figures is not based on empirical findings, but is contrived. At most, some sophisticated guessing is involved in places.

In Figure 4a, the words *vader* and *father* again (cf. Figure 3a) share all their meaning elements. The same is the case with the words *moeder* and its translation *mother*. In addition, the words *vader* and *moeder* have three elements in common. But because both *vader* and *moeder* share all their conceptual elements with their respective translations, *vader* also has the same three elements in common with *mother*, and, conversely, *father* has the same three in common with *moeder*.

Figure 4b depicts a situation in which Figure 3b is expanded. The words within both pairs of translations share four meaning elements. The semantically related words share two meaning elements both within and between languages.

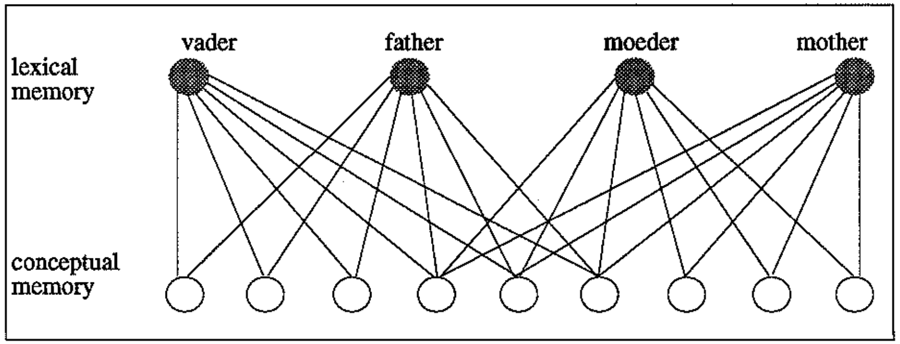

Figure 4a. **Distributed conceptual representations in memory. Translations share all meaning elements. Semantically related words share a few, both within and between languages.**

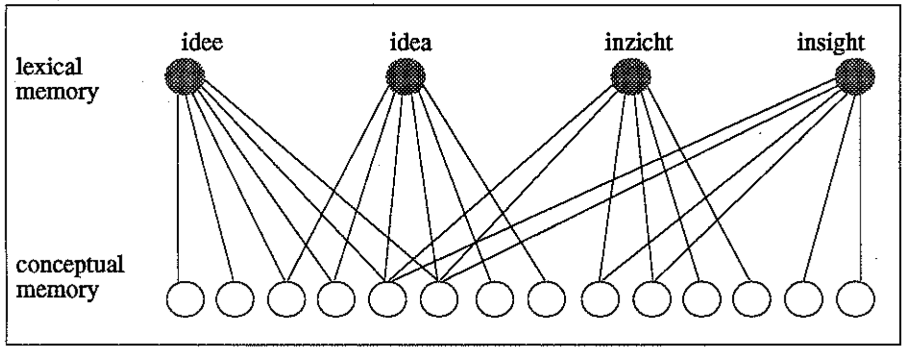

Figure 4b. **Distributed conceptual representations in memory. Translations differ in meaning. Semantically related words share a few meaning elements, both within and between languages.**

The representations depicted in Figures 4a and 4b could account for a very common translation error, namely, the production of a word semantically related to the stimulus word, but in the language of the translation. My work on word translation has abundantly provided me with such errors. Here are a few examples (the first word is the Dutch stimulus word to be translated. Its translation is given in between brackets. The last word is the response actually given): kaars (candle)—flame; speld (pin)—needle; cirkel (circle)—square; boomgaard (orchard)—vineyard; handdoek (towel)—blanket; bruid (bride)—wedding; bliksem (lightning)—thunder; aardappel (potato)—carrot; plafond (ceiling)—roof. These errors can be understood as arising from the activation of the lexical node of the response word (e.g., square) via the nodes representing the meaning elements it shares with the stimulus word (cirkel). The more elements shared between a stimulus word and a word in the other language that is not its translation, the larger the activation in the lexical node for the latter will be, and hence the larger the chance that the associated error will occur.

Translation recognition

In my work on word translation I explored a new version of this task, which I labeled 'translation recognition.' Bilingual subjects performing this task are presented with pairs of words, each consisting of a word in their first language and one in their second language. The task is to decide whether or not the words within each pair are translations of each other. In terms of Figure 1, translation recognition may involve the tracing of the links (a spread of activation) departing from the lexical nodes of the two presented words. Both the direct links at the lexical level and, if present, the indirect ones via the conceptual level, will be traced. If an intersection of activation occurs anywhere, a *yes* response is emitted. If not, the subject responds *no*. On negative trials (requiring a *no* response) the subjects in my experiment were always presented words that were not related in any obvious sense (phonologically, orthographically, or semantically). The searches starting from the lexical nodes of the presented words will thus intersect nowhere, and a *no* response will be given. According to the distributed view of representation discussed here, translation recognition again entails activation spreading from the lexical nodes of the words involved and the detection of intersecting activation waves, but now the routes involved are the links between the lexical nodes and the representations of the individual meaning elements at the conceptual level.

From the present conception of representation a prediction can be derived concerning an experimental condition I have not tested yet, namely one in which the negative trials consist of semantically related words (e.g., the word pair *vader-mother*). Because an intersection of activation will occur on one or more links (three in the example), the subjects will be biased towards a *yes* response. Consequently, the rejection of these words as a pair of translation equivalents should take relatively long, and relatively many errors should occur.

More tasks and more effects are hidden in Figures 4a and 4b, for instance, semantic priming effects within and between languages. They are the topic of the following section.

Semantic priming within and between languages

A robust effect in monolingual investigations is that words in, for instance, a lexical decision experiment are responded to faster when they follow a semantically related word (e.g., context stimulus or 'prime': *love;* test stimulus: *friendship*) than when preceded by a semantically unrelated word or some neutral context stimulus (Meyer & Schvaneveldt, 1971; see Neely, 1991, for a review). This 'semantic-priming' effect is often, again, attributed to activation spreading between memory nodes, for example, in a memory system of the type illustrated in Figure 1: If a word is presented that corresponds to one of the representations preactivated through activation spreading from the representation of an earlier prime, it is recognized, and hence responded to, relatively fast. In this view, recognition is effectuated as soon as the activation in the test word's lexical node exceeds a critical threshold value. In Figure 1, when *vader* is the prime and *moeder* the test word, responding to the latter comes about relatively fast because the lexical node for *moeder* has received preactivation from the lexical node of *vader* via the conceptual nodes for *vader* and *moeder*.

Priming effects of words presented in one language on semantically related words in a second language can be explained in the same way. If translation equivalents share a conceptual representation, when *vader* is presented as prime the lexical node for the test word *mother* should also receive preactivation, via the conceptual nodes *vader/father* and *moeder/mother*. In fact, under the scheme of the top half of Figure 1 (and given equally strong T2a and T2b connections), the lexical node *mother* should receive the same amount of preactivation as the node for *moeder*, and the between-language priming effect should thus be as large as the within-language effect. The bottom half of Figure 1 depicts a situation wherein within-language semantic priming should occur, but no between-language priming.

A number of studies have shown that a semantic-priming effect can occur between languages (Chen & Ng, 1989; de Groot & Nas, 1991; Jin & Fischler, 1987; Kerkman, 1984; Kirsner, Smith, Lockhart, King & Jain, 1984; Meyer & Ruddy, 1974; Schwanenflugel & Rey, 1986; Tzelgov & Henik, 1989). Some of them (Chen & Ng, 1989; Meyer & Ruddy, 1974; Schwanenflugel & Rey, 1986; Tzelgov & Henik, 1989) suggested that the interlingual effect is as large as the intralingual effect. Two studies indicated that the interlingual effect may be word-type specific. Jin and Fischler (1987) observed a semantic-priming effect across languages for concrete words but not for abstract words. Under conditions in which the prime was degraded so that it could not be identified by the subjects, Gerard Nas and I (de Groot & Nas, 1991) obtained a between-language effect for cognates but not for noncognates. This is in fact why in the bottom part of Figure 1, which shows the representation of noncognates, separate conceptual nodes for translation equivalents, and only within-language connections between semantically-related words at the conceptual representational level, are depicted. The assumption of separate representations directly follows from the absence of an interlingual effect for noncognates. Had there been a shared representation, such an effect should have occurred.

Figures 4a and 4b show how semantic priming, both within and between languages, comes about when conceptual representations are distributed across a number of nodes (see also Taylor & Taylor, 1990, and, for monolingual semantic priming, Masson, 1991). By the time the test word is presented, activation has accumulated in its lexical node via the representations of the meaning elements it has in common with the prime. The more elements the prime and test words share, the larger the amount of preactivation in the test word's lexical node, and the larger the priming effect should be. Thus, the effect should be larger in Figure 4a than in Figure 4b (three and two elements shared by the semantically related words, respectively). Differences in the number of conceptual elements that the prime and the test word have in common could explain the finding in the monolingual literature on semantic priming (e.g., de Groot, Thomassen & Hudson, 1982; Fischler & Goodman, 1978) that the size of the priming effect depends upon the 'strength' of the semantic relation between the prime and the test words. Of course, a representational structure of the kind depicted in Figure 1 could also explain such an effect, namely by assuming that the links between the conceptual representations of semantically related words differ in strength, the strength of each individual link reflecting the strength of the semantic relation between the two words represented in the nodes on both ends of the link. The stronger the link, the more activation it will transmit, and the larger the effect will be.

When one compares one of the accounts of word translation in terms of the representations depicted in Figure 1 (Route T2), on the one hand, with that in terms of the representations in Figures 4a and 4b, on the other, one might want to argue that they are in fact conceptually the same. But the two conceptions of semantic priming differ essentially. Given the representations in Figure 1, for semantic priming to arise, activation in one conceptual representation (of the prime word) has to traverse a link from this representation to another (of the test stimulus) in conceptual memory. According to the view depicted in Figures 4a and 4b, no such links between conceptual representations have to be traversed. They may even not exist. All links responsible for the effect directly connect nodes in lexical memory with nodes in conceptual memory. The priming effect is due to the fact that the prime, by activating its distributed conceptual representation, at the same time activates part of the conceptual representation of the test word.

Representations as in Figures 4a and 4b could also readily account for differences in the size of the semantic-priming effect within and between languages. I will not consider the hypothetical situation where larger between- than within-language effects are obtained. I do not know of any study in which such a finding was obtained, and it seems intuitively implausible. But priming effects may be smaller between than within languages. Models with non-distributed conceptual representations could explain such results in two ways: 1) There are between-language connections between the representations of semantically related words in conceptual memory for all words for which the corresponding within-language connections exist, but they are weaker than the latter. This option demands the existence of language-specific conceptual nodes. (In the case of shared representations, the within- and between-language connections between two nodes in conceptual memory would in fact be one and the same connection. It is hard to see how this one connection could be strong and weak at the same time.) Thus, for example, there would exist a conceptual node for *liefde* and one for its translation *love;* one for *vriendschap* and one for its translation *friendship*. Additionally, there would be relatively strong links between the nodes for *liefde* and *vriendschap,* and between those for *love* and *friendship,* and there would be weaker links between the nodes for *liefde* and *friendship*, and between those for *love* and *vriendschap*. 2) The between-language connections between the representations of semantically related words in conceptual memory are as strong as the corresponding within-language connections, but they do not exist for all of the semantically related words that are connected *within* a language. The situation depicted in Figure 1 is one way to instantiate this second option. Here a subset of the translation equivalents in the bilingual lexicon (cognates) shares a conceptual representation. The reason the between-language connections are as strong as the corresponding within-language connections is that they are in fact the same connections. The remaining words (noncognates) are represented in language-specific conceptual nodes that are only connected to conceptual representations of semantically related words of the same language. But also compatible with Option 2 would be a situation in which *all* translation equivalents would be represented in language-specific nodes. So there would be, for instance, two word quartets, *vader-father-moeder-mother* (cognate translations), and *oom-uncle-tante-aunt* (noncognate translations). Each of these eight words would individually be represented in a conceptual node. All intralingual connections between the conceptual representations of semantically related words, irrespective of the cognate status of the words involved, could be equally strong (*vader-moeder*; *father-mother*; *oom-tante*, and *uncle-aunt*). The interlingual connections between the representations of semantically related cognates (*vader-mother and father-moeder*) could also be this strong, but no connections would exist between *oom* and *aunt*, or between *uncle* and *tante* (noncognates). Options 1 and 2 could be distinguished by item analyses on the data, because the interlingual priming effect should be significant by items if Option 1 were true (all or the majority of the interlingually semantically related word pairs would show the effect), but not if Option 2 were true (only a subset of these word pairs would show the effect).

These solutions are relatively complex and may even appear contrived. They are certainly more complex than the one the distributed view has to offer: It does not seem far-fetched to assume that semantically related words of the same language often share more meaning elements than semantically related words of different languages. The larger the overlap, the more activation will accumulate in the lexical node of the test word, and the larger the priming effect will be. Hence, the effect will be larger within a language than between languages.

In this framework, when for particular types of words (abstract words, Jin & Fischler, 1987; noncognates, de Groot & Nas, 1991) a within-language but no between-language priming effect is obtained, one is not compelled to conclude that the conceptual representations of those words are strictly separated by language, as one is when conceptual representations are regarded as indivisible entities. The translations of such words may still have a large part of their conceptual representations in common, but these words would not share any of their conceptual elements with semantically related words in the other language.

Word association

Word association has also been used as a tool to investigate bilingual memory (Kolers, 1963; Taylor, 1971; 1976). There are two common versions of this task: discrete word association and continued word association. In the former the associative response to a stimulus word has to consist of a single word that is the first word that comes to the subject's mind when reading or hearing the stimulus word. In the latter version, the subject generates as many word associates to the stimulus word as possible within a prespecified amount of time (often 30 or 60 seconds). In bilingual word-association studies stimulus words are typically presented in one or both of the bilingual's two languages, and responses have to be given either in the language of the stimulus word, or in the other language. The issue at stake is to what extent the responses in the various experimental conditions are or are not the same (responses that are translations of those given in other conditions are considered 'same' responses). Same responses are regarded as supporting the view of conceptual representations being shared between languages. Different responses are seen as evidencing language-specific conceptual representations.

Kolers (1963) collected discrete word associations within and between languages. His subjects all had English as their second language, and German, Spanish, or Thai as their native language. Each individual subject produced associations in each of the four within- and between-language conditions. His main finding was that within all three groups of bilinguals over half of all responses in the cross-language conditions were unique, that is, not the same as or a translation of the response word this particular subject gave in either one of the two within-language conditions. He concluded that 'experiences and memories of various kinds are not stored in common in some supralinguistic form but are tagged and stored separately in the language S used to define the experience to himself' (Kolers, 1963, p. 300). This conclusion may be too strong, given the fact that at least a number of responses were shared between languages: On average, just over 20% of the responses of an individual subject were the same as or a translation of those she or he produced in all remaining conditions (e.g., *king-queen; king-reina; rey-reina; rey-queen*. Examples are taken from Kolers, 1963). Furthermore, about 30% of the interlingual responses were the same as or a translation of the response word this subject gave in either the native or nonnative intralingual condition (e.g., *boy-girl; boy-nina; muchacho-hombre; muchacho-trousers*. In this example the subject's response in the nonnative-to-native condition *[boy-nina]* was the same as his response in the nonnative-to-nonnative condition *[boy-girl]*). Note that in the examples above the response words are always in some sense semantically related to the stimulus words. This reflects a fact that always immediately strikes any student of word association: Although the task instructions never explicitly demand this, by far the majority of word association responses indeed are words semantically related to the corresponding stimulus words.

A further interesting finding of Kolers (1963) is that concrete words more often generated the same responses within and across languages than abstract words and emotion words did. The former result was also obtained by Taylor (1976), who tested French-English bilinguals in intra- and interlingual continued word association. Additionally, she observed that stimulus words with cognate translations more often gave rise to the same response words in the intra- and interlingual conditions than noncognates did.

Before drawing conclusions on the basis of these data on the organization of bilingual knowledge in memory, one would first want to know about the chances that a subject will respond with the same word when associating to a word twice within the same language. But for the time being the data suggest, first, that words and their translations in bilingual memory neither fully share their conceptual representations nor are represented in a totally segregated way, and, second, that the degree of separation between languages varies with word type.

Like word translation and semantic priming discussed before, the word association task, both within and across languages, can also be detected in Figures 4a and 4b. In within-language association, viewed in terms of the memory structures suggested here, the same paths are involved as in within-language semantic priming (recall that the response words in word association are typically semantically related to the stimulus words). In between-language association either a laborious process may take place, or a simpler one. Kolers (1963) assumed that in cross-language word association, bilinguals either first translate the stimulus word and then associate to the translation, or they first associate to the stimulus word in the language of the stimulus and subsequently translate the association. Both of these indirect routes are visible in Figures 4a and 4b (e.g., from *vader* via conceptual memory to *father*, and from there, again via conceptual memory, to *mother*; or from *vader* to *moeder* to *mother*, again both via conceptual memory), but a direct route (as direct as the within-language word association route in terms of this type of representation would be) can also be discerned, from *vader* into conceptual memory and from there straight to the lexical node for *mother* (cf. interlingual semantic priming).

On any trial a number of lexical nodes for words that would all constitute appropriate responses will be activated. The one activated most will generally determine the response. So if, in an intralingual condition, after presentation of the stimulus *father* the lexical node for *mother* receives more activation than the nodes for any of the other same-language words plausibly being activated (for instance, *son* and *child*), the corresponding word *mother* will be produced as response. If in an interlingual condition following the presentation of this same stimulus word *father* the lexical node for *moeder* is activated more than any of the other lexical nodes of words in the target language, this node will determine the response. In this situation, the within- and between-language responses will thus be the same. But if the lexical node *zoon (son)* is activated more than is the node for *moeder*, for instance, because the conceptual representation of *zoon* shares more elements with that of *father* than the conceptual representation of *moeder* does, the intra- and interlingual conditions will give rise to different responses. The association data suggest that for some types of words (concrete words; cognates) the maximum activation in the intra- and interlingual conditions relatively often (as compared to abstract words, noncognates, and emotion words) occurs in the lexical nodes of translation equivalents.

The tasks discussed so far may be classified into three groups: production tasks, priming tasks, and relation-assessment tasks. In the *production* tasks a stimulus word is presented from which the subject has to generate a particular type of response. All such tasks implicit in Figures 4a and 4b have been explored so far. These tasks were: word translation (in Figure 4a, from *vader* to *father* and vice versa, and from *moeder* to *mother* and vice versa); within-language word association (from *vader* to *moeder* and vice versa, and from *father* to *mother* and vice versa); and between-language word association (from *vader* to *mother* and vice versa, and from *father* to *moeder* and vice versa).

In the *priming* tasks the subjects have to respond to target stimuli preceded by a prime. The required response could be lexical decision, but other responses may be requested instead (e.g., pronouncing the targets, or performing some semantic classification of them). The influence of the prime on target processing is assessed. Unlike in the relation-assessment tasks to be discussed below, the prime may be ignored by the subjects. The tasks of this type hidden in Figures 4a and 4b are: intralingual semantic priming (prime: *vader*, target: *moeder*, and vice versa, and prime: *father*, target: *mother*, and vice versa), interlingual semantic priming (prime: *vader*, target: *mother*, and vice versa, and prime: *father*, target: *moeder*, and vice versa), and interlingual repetition priming (or 'translation' priming; prime: *vader*, target: *father*, and vice versa, and prime: *moeder*, target: *mother*, and vice versa). The first two of these have already been discussed, but translation priming has been ignored so far. It is the topic of the next section. A characterization of the third group of tasks, the *relation-assessment* tasks, is postponed until later, when a few examples of this class of tasks will be discussed.

Translation priming

Translation priming or between-language repetition priming has been looked at in a large number of studies (Altarriba, 1992; Chen & Ng, 1989; Cristoffanini, Kirsner, & Milech, 1986; de Groot & Nas, 1991; Gerard & Scarborough, 1989; Jin & Fischler, 1987; Kerkman, 1984; Kirsner, Brown, Abrol, Chadha, & Sharma, 1980; Kirsner et al., 1984; Scarborough, Gerard & Cortese, 1984). In all but four of them the 'classical' interlingual repetition-priming paradigm has been used. In this paradigm, the inter-stimulus-interval between a word and its translation is typically long, several minutes or more, and the subjects produce some response to both the word and its translation. In the four remaining investigations (Altarriba; Chen & Ng; de Groot & Nas; Jin & Fischler), as in studies on semantic priming, a word and its translation (or some other test stimulus) followed one another immediately (across the studies, the stimulus-onset-asynchrony between prime and test stimulus varied between 60 ms and 1000 ms), and the subjects only responded to the latter. In all four of these studies translation priming occurred. The effect occurred not only when the prime was clearly visible (Altarriba, 1992; Chen & Ng, 1989; Jin & Fischler, 1987), but also when it was masked so that it could not be identified by the subjects (de Groot & Nas, 1991). We thought masking the prime to be relevant because, when the experimental task is lexical decision (true for all four studies) and when both the prime word and the test stimulus are clearly visible, a post-lexical integration process may also cause a priming effect. This post-lexical process searches for any relation, for instance, a translation relation, between prime and test stimulus. If it finds one before the subject executes his or her response to the test stimulus, it speeds up this response (see de Groot & Nas, 1991, for details). Whenever the primes are not masked it is thus not clear to what extent the effect may be attributed to the actual priming process.

In de Groot and Nas (1991, Experiments 3 and 4) the cognate status of the translation equivalents was varied. Considering the masked-prime condition only, the effect was always larger for cognates than for noncognates (in one condition the difference in effect size was substantial: 53 ms), although statistically the effect was always equally large for the two types of words. The language combination studied by Jin and Fischler (1987) was Korean-English; that studied by Chen and Ng (1989) was Chinese-English. Unlike in the English script, the units in both Chinese and Korean script are characters. Consequently, Korean-English and Chinese-English translations will always be orthographically dissimilar. They will also generally be distinct phonologically (except that words imported from English into Chinese and Korean or vice versa may retain aspects of the pronunciation of the imported words). In short, the stimulus materials of Jin and Fischler and of Chen and Ng consisted of noncognates. Altarriba's (1992) subjects were Spanish-English bilinguals. The languages involved are both alphabetic, but belong to different language families (Romance and Germanic, respectively). Therefore, her translations probably also consisted primarily of noncognates. Despite the use of noncognates as stimulus materials, translation priming was obtained in all three studies.

In the studies using the classical paradigm, the interlingual effect is less robust, but there is a pattern: Translation priming occurs for cognates (Cristoffanini et al., 1984; Gerard & Scarborough, 1989; Kerkman, 1984), but not for noncognates (Kirsner et al., 1980; Kirsner et al., 1984; Scarborough et al., 1984). However, there are grounds to doubt that the effect under the conditions of these experiments is attributable to spreading activation in bilingual lexical memory, which is our concern here. Instead, it may be an episodic effect (see de Groot & Nas, 1991, for a discussion).

If the representations in Figure 1 are the building blocks of bilingual memory, translation priming for cognates (at least in studies where the non-classical paradigm, the one modeled on semantic-priming studies, is used) could come about through activation spreading directly, via Route T1, or indirectly, via Route T2, from the lexical representation of the prime word to that of the test word, preactivating it prior to its presentation (cf. word translation). For noncognates preactivation could only come about via Route T1, because no indirect connections via conceptual memory exist. Recall that we (de Groot & Nas, 1991) assumed language-specific conceptual representations for noncognate translations because for noncognates no interlingual semantic-priming effect was obtained. However, a translation-priming effect *did* occur. The combination of these two findings forced us to conclude that direct links exist between the lexical nodes of translation equivalents. If indeed no indirect connections between these translations via conceptual memory exist, how else could translation priming for noncognates be explained?

Unlike the view of representation illustrated in Figure 1, the present view does not require the conclusion that direct (T1) connections exist between the lexical representations of translations. They may exist (indeed others have proposed their existence for different reasons; see the section on word translation), but they do not have to. The data summarized above can no longer be regarded as conclusive about this. If the conceptual representation is divided over a number of different nodes, it is perfectly plausible that for a particular type of word (presently noncognates) translation priming occurs, and does so via conceptual memory, whereas at the same time no interlingual semantic-priming effect for this type of words comes about. What would be required is (at least partially) overlapping conceptual representations for a pair of noncognate translations, while at the same time none of the nodes representing the various meaning elements in these conceptual representations is linked to the lexical node of the relevant target word in an interlingual semantic-priming condition.

Relation assessment

One of the tasks discussed so far, translation recognition, may be considered an instance of a class of tasks in which the subjects have to decide whether or not a particular relation between two stimuli exists. These tasks necessarily involve the processing of both stimuli on a trial. In this respect they differ from the above priming tasks (excluding the 'classical' repetition-priming studies), in which the subjects may ignore the first stimulus within each pair of stimuli. Other instances of this group of tasks implicit in Figures 4a and 4b would be intra- and interlingual semantic-relation-assessment tasks, which would require subjects to categorize word pairs according to the presence or absence of *any* semantic relation between the words in these pairs. If such relation is detected, as with the pairs *vader-moeder* (intralingually) and *vader-mother* (interlingually), the subject should respond *yes*. If not (*vader-boom*, or *vader-tree*), *no* should be the response.

Analogous to the conception of translation recognition, semantic-relation assessment may be conceived of as involving activation spreading from the lexical nodes of the two presented words. If an intersection occurs, a *yes* response can be emitted. If not, a *no* response may be executed. I do not know of any study in which it is the subjects' task to categorize the presented word pairs on the presence or absence of *any* semantic relation between the words of a pair, but this hypothetical task is strongly reminiscent of the more specific 'semantic-verification' task that has been used in a very large number of studies (e.g., Collins & Quillian, 1969; Smith, Shoben, & Rips, 1974). In semantic verification as well a relation between the two words on a trial has to be discovered, but the relation to be detected has to be of a specific kind. Other than in the above task, if the words on a trial are semantically related, but not in the prespecified way, such a trial demands a *no* response. In one study of this type (Caramazza & Brones, 1980), semantic verification was investigated both intra- and interlingually.

Caramazza and Brones presented word pairs on a screen, the first word referring to a semantic category, and the second to an instance of this or another category. Subjects had to press one key if the second word belonged to the category referred to by the first, and to press another key if such was not the case. Three categories and six instances of each of them constituted the experimental materials. The categories were 'furniture,' 'fruit', and 'vegetables.' Hence, all stimuli were concrete words. The category and instance names were in the same or in different (English and Spanish) languages. The finding most relevant here was that response time was not influenced by whether or not the names of category and instance were in the same language. Two robust findings in semantic verification studies were replicated by Caramazza and Brones in their cross-language condition: (1) correct *yes* responses took less time when the instance was typical of the corresponding category *(fruit-apple)* than when it was atypical *(fruit-melon)*, and (2) correct *no* responses took longer when the instance was drawn from a category semantically related to the category mentioned on the trial *(fruit-carrot)* than when drawn from a semantically unrelated category *(fruit-chair)*.

These findings can readily be understood in terms of the distributed conceptual representations proposed here, by assuming that the critical variable in the decision process is the number of conceptual elements a category shares with the instance presented on the same trial (cf. the interpretation of Smith et al., 1974, in terms of the number of shared features). Three specific assumptions need to be made: (1) A typical instance shares more conceptual elements with its category than an atypical instance. (2) Not only does an instance share conceptual elements with the category it belongs to, but it also shares some with a related category. (3) An instance and a semantically unrelated category do not have any of their conceptual elements in common. When there are many common elements (typical instance) and, hence, a large amount of activation at the intersection of the activation waves spreading out from the two presented words, the subject assumes the instance belongs to the specified category, and immediately responds *yes*. When there are no shared elements (unrelated non-instance) and, hence, no area of intersecting activation in conceptual memory, the subject assumes the instance does not belong to the specified category and responds *no* relatively fast. In the case of a few shared elements (atypical instance; related non-instance) and, therefore, some activation at the intersection, the subject has to be on guard, because either *yes* or *no* may be the correct response. He or she must somehow evaluate the links of the intersection, a process taking additional time. Consequently, the response times are relatively long on these trials. The decision process in the semantic-verification task is thus more complex than in the above general semantic-relation-assessment task, because in the latter the evaluation stage is redundant: Any intersection of activation indicates a relation, so whenever an intersection is detected, a *yes* response is appropriate (even on trials of the related non-instance type).

The fact that the response pattern in Caramazza and Brones' study was independent of the language of the stimulus materials suggests that the English and Spanish words for the (concrete) categories and instances used in their study shared the same set of conceptual elements in these bilinguals' memories.

A second bilingual investigation that belongs in this section is an unpublished study by Colletta, reported by McCormack (1977). It resembles that of Caramazza and Brones (1980) in that not *any* but a specific type of relation had to be searched for. Colletta's subjects were English-French bilinguals. They were presented with word pairs and had to decide for each individual pair whether or not it consisted of synonyms. The words within a pair were presented either in the same language or in different languages. Response times in the intra- and interlingual presentation conditions were equally long. This finding was seen as support for the view that translations share a representation in bilingual memory. In the present terms it again suggests that the corresponding words in the two languages share the same set of conceptual elements.

Integration or segregation?

Many an opening paragraph of writings on bilingual lexical organization states that the lexical knowledge of the bilingual may be represented in two language-specific memory stores, one for each of this bilingual's languages, or may instead be integrated in a single language-independent store. A tenet of the foregoing has been that the truth may lie somewhere in between these two extreme positions. Some words may have all of their conceptual representation, others relatively little or maybe even nothing in common with their closest translation. Another suggestion made in this chapter is that individual words may or may not share part of their conceptual representation with a semantically related word in the other language (and in the same language, but that is of less interest here). The data reported in this chapter suggest that the emerging representational form is likely to depend on word type (Is the word abstract or concrete? Does it evoke particular emotions or is it emotionally neutral? Is it a cognate or a noncognate?). But the degree of overlap between the meaning of a particular word and that of its closest translation may vary with word type, so it is possible that ultimately not word type *per se*, but the extent to which the meanings of the translations overlap, is the critical factor that determines how the two are stored in memory.

At various points in this chapter word concreteness was mentioned as a determinant of bilingual task performance. If amount of meaning overlap indeed underlies the effects of this variable (meaning overlap determining the amount of sharing between the conceptual representations, and the latter, in turn, determining the effects), concrete words and their translations must be more similar in meaning than abstract words and their translations. Although empirical data will have to be collected to substantiate it, the view that the meanings of concrete words are more similar across languages than those of abstract words is intuitively very plausible. The function of the entities referred to by concrete words will generally be the same in different language communities. Wherever we come across them, chairs are to sit on, and apples to eat. The appearance of these entities will also generally be the same across different language communities. That of man-made objects like chairs will to a large extent be imposed on them by their function, and hence be similar across different communities. That of natural objects like apples will generally be the same everywhere by virtue of the fact that they are natural categories. The end-product of learning a concrete word will thus be a representation of which the content varies relatively little across languages. Abstract words have no external referents that could be looked at, handled, utilized, and thus guarantee similarity of the content of the developing representations across languages. Their meanings have to be acquired by looking up these words' definitions (but see the next section) in a dictionary (or asking others to provide them), and, more importantly, by deducing them from the various contexts in which these words are used. To the extent that these contexts differ between languages (cultures), the meanings of these words will also differ. In sum, there are good grounds to assume that concrete words and their translations have very similar meanings, whereas abstract words and their translations have meanings that differ more substantially. Consequently, the chances that abstract words are represented language-specifically are larger than for concrete words.

For one group of abstract words an abundance of literature exists which bears on the present issue: Many ethnographic studies have been concerned with the meaning of particular emotion words in the community under investigation. Russell (in press) reviews the relevant literature and provides a wealth of examples suggesting that the reference of these words often differs between languages (cultures). A first indication for this is that languages differ considerably in terms of the number of words they possess to categorize emotions. The number of emotion words in different languages may vary between over two thousand at one extreme (in English, although only a minority of these may be in the vocabularies of individual speakers of English) and only seven at the other (in the Chewong language; Russell, in press). This may be taken to indicate that there are large cross-cultural differences in the extent to which people experience emotions (the fewer emotions, the fewer emotion words, or/and vice versa), but it may also indicate that the meaning of an emotion word and that of its closest translation in another language differs (of course, both may be the case). For instance, each of the emotion words in languages that contain only a few of them may cover more than the corresponding words in languages with a richer emotion vocabulary. But even when two languages have an equally large emotion vocabulary, the reference of corresponding words in the two languages may differ. Given two languages L1 and L2, some of the emotion words in L1 may have a broader, others a more narrow reference than the corresponding words in L2. Also, L1 words may exist for concepts that cannot be expressed in a single word in L2 or that do not exist as, or cannot even be conceived as, concepts in L2, and vice versa (the reader is referred to Russell, in press, for a thorough documentation of all these situations). All these words are likely candidates for language-specific representation in the memory of a bilingual whose two languages are L1 and L2. Kolers' (1963) word-association data discussed earlier support the present view that emotion words are relatively often represented language-specifically.

Some of the studies mentioned in this chapter suggested that, besides word concreteness and the emotional content of words, cognate status of the translation equivalents is yet another determinant of bilingual performance. If, again, the degree of meaning overlap between the translations is the critical factor underlying the observed effects (with representational form as mediator), cognate translations must have more similar meanings than noncognate translations. A reason for this could be the differential origin of cognate and noncognate translations: Cognate translations will generally derive from the same root in a common parent language. If they have both preserved the meaning of this root, or at least a large part of it, over time, they will have ended up having (about) the same meaning. In contrast, noncognate translations will generally not derive from the same root, and there is thus a relatively large chance that their meanings will differ more between the languages. But it may also be that not (or not only) differences in the degree of meaning similarity but (also) in perceptual similarity cause cognate translations to come to share more of their conceptual representations than noncognate translations: L2-learners, noticing the orthographic and phonological similarity between a cognate word and its translation, may simply assume the two have the same or about the same meaning, and thus conveniently link the new L2-word onto the conceptual representation of the corresponding L1-word. An interesting consequence of this may be that orthographic and phonological similarity of translations thus blind the learner to differences, if any, between the meanings of these two words.

There are other factors that may affect the way translations are stored in bilingual memory, but that I will only touch upon here. One is the circumstances under which the languages in question are acquired. Ervin and Osgood (1954) suggest that bilinguals who learn their two languages in different environments ('coordinate' bilinguals) develop a memory structure with separate representations for word translations in the two languages, whereas those who learn their languages by using them interchangeably ('compound' bilinguals) develop a memory structure with representations that are shared by the two translations. For instance, the common practice in foreign language classrooms where an L2 word is taught by directly associating it with its translation in the native language is one way to create compound bilinguals (see Keatley, 1992, for a longer discussion of this and related distinctions). Yet another critical factor may be whether or not the person's two languages belong to the same family: The chances that translations share a representation in memory (or share a relatively large part of their representations) may be larger when the languages of the bilingual are related than when unrelated. But here, again, the degree of meaning similarity between the translations may ultimately be the critical factor. Translations of words belonging to related languages may be more similar in meaning than those of words in unrelated languages. Not the presence or absence of a relationship between the languages *per se*, but this meaning similarity at the level of individual words, may determine how they will eventually be represented in memory.

The contents of the conceptual representations

The starting-point of the view on bilingual lexical memory set forth in this chapter was that a conceptual representation is composed of a number of conceptual elements. Until now nothing has been said about the nature of these elements. In this last section some of the relevant literature will be discussed.

According to what is known as the 'classical' or 'traditional' view of concept representation, concepts are represented by a fixed list of features that together *define* the concept (e.g., Katz, 1972, but dating back much longer, to the Greek philosopher Aristotle); that is, the features are individually necessary and jointly sufficient for membership of the category (I am using the terms *concept* and *category* interchangeably here, as is done more often). Assuming that concepts can be defined implies that the boundaries between concepts are clear-cut and stable. For instance, it should be clear where exactly cups turn into bowls, and where bowls turn into plates. Another implication of the classical view is that members of a category have equal status, that is, each member should be as good a member of the category as any other member. Yet another is that a given concept does not vary within the same individual or across individuals.

Many empirical findings and thought-experiments have cast doubt on this view of concept representation. I will mention just a few here. Wittgenstein (1953), to name an illustrious opponent, took the concept 'game' as an example with which he challenged it. He argued that this concept (and most others) cannot be captured in terms of a set of features that holds for all instances of this category. Rosch (e.g., 1973) collected experimental evidence suggesting that individual members of a category do not have equal status. Instead, many categories appear to have a 'graded' structure, with some members being more typical of the category than other members (a chair is a more typical

instance of the category 'furniture' than a clock is). Typical members are those that can be captured in a set of 'prototypical' or 'characteristic' features that are listed for many (not *all*) members of this category, whereas atypical members can be described with a list that contains less of these features common to the category, but instead contains relatively many features not shared by most of the other members.

Experimental studies showing an effect of context on categorization constitute yet another serious challenge to the classical view. In an early study Labov (1973) collected data suggesting that concept boundaries are not clear-cut and static, but vary with context. His subjects had to categorize pictures of objects like cups, bowls and plates by naming the depicted object. Prior to naming them, they were instructed to, for instance, imagine someone holding the object and drinking coffee from it or, in a second condition, to imagine the object filled with mashed potatoes and sitting on the dinner table. It turned out that the objects were classified differently in different contexts: One and the same object was classified relatively often as a cup in the first of the above contexts, and relatively often as a bowl in the second. In related work, Barsalou (1987; Barsalou & Medin, 1986) provides a wealth of experimental data indicating that concepts vary with context, both 'long-term' context (people's experiences) and current context (e.g., linguistic context or point of view). For instance, two individuals' concept of 'bird' may differ because of different experiences of these individuals with birds, but the concept of one and the same individual may also differ at two different points in time because of this individual's new experiences with birds in the intervening period. In other words, people's representation of categories reflects their experiences (Barsalou & Medin, 1986). That concepts also vary with *current* context can be concluded from a study by Barsalou and Sewell (in Barsalou, 1987). Subjects judged instance typicality from one of several international points of view (for instance, from the American and Chinese points of view). Groups of subjects (sampled from the same population) taking different points of view produced different graded structures for the same category (for instance, *robin* and *eagle* were judged to be typical instances of the category 'bird' from the American point of view, whereas *swan* and *peacock* were considered typical from the Chinese point of view). Discussing the variability of concepts, Aitchison (1987, p. 40) uses some lively metaphors. She likens concepts to elusive butterflies and slippery fish: 'Word meanings cannot be pinned down, as if they were dead insects. Instead, they flutter around elusively like live butterflies. Or perhaps they should be likened to fish which slither out of one's grasp.'

All these studies thus indicate that it is not a fixed set of defining features that conceptual representations typically consist of (although *some* concepts may be represented that way). The studies showing an effect of current context on the content of concepts (Barsalou and Medin's point of view experiment, but also a study by Roth and Shoben, 1983, showing an effect of current linguistic context) suggest that people *construct* representations that suit the context (Barsalou & Medin, 1986). The clearest demonstration of this is by Barsalou (1983). He showed that people often construct new categories to achieve a current goal. These 'ad hoc' categories differ from common categories in that they are not well established in memory. Examples from his study are 'ways to make friends', 'things that could fall on your head,' and 'ways to escape being killed by the Mafia'. It could be argued that this study does not bear on the representation and processing of *common* concepts, but the data suggest otherwise: Ad hoc categories possess the same graded structures as common categories.

In short, conceptual representations appear to be constructed when needed. Does this imply that *all* of the concept is built up on-line, in other words, that there is no permanent representation in memory to be accessed as a whole each time the corresponding word is encountered, irrespective of context? If so, the view of representation set forth in this chapter would run into trouble, because it assumes the existence of such static part in the concept representations. Fortunately, it seems that a relatively stable *core* representation may still be assumed. These cores are generally not definitional (because word definitions exist for few words) but experientially-based (Barsalou & Medin, 1986). Barsalou (1982) distinguishes between context-independent and context-dependent properties (features) in concepts. Context-independent properties are activated each time the corresponding word is encountered (e.g., the property 'smells unpleasantly' when encountering the word *skunk*). Context-dependent properties are activated only by a relevant context in which the word occurs (e.g., the property 'floats' of the concept 'basketball' in the following sentence: *Chris used a basketball as a life preserver when the boat sank*; examples are taken from Barsalou, 1982). A concept's set of context-independent properties may constitute a relatively stable core representation in memory. *Relatively* stable, not just stable, because these cores are experientially-based. With new experiences the cores may change somewhat. But many of their properties will be immune to changes: Most birds will go on flying forever and chairs will always be for sitting. Furthermore, many of these properties will hold across languages. They are thus plausible candidates for the language-independent conceptual elements I have assumed in the preceding sections.

An interesting possibility to ponder is that the size of these cores varies with word concreteness (and maybe with other word characteristics as well), the cores of abstract words containing fewer elements than those of concrete words. One reason to consider this is the casual observation that, when asked to define a word, our response often is particularly clumsy in the case of abstract words. We often do not fare well with concrete words either (which is not surprising, given the fact that for most words definitions do not exist), but at least we can come up with some information on the associated concept (e.g., its characteristic properties; its function, if any; the superordinate; a number of subordinates). In the case of abstract words it seems that we can often only think of a number of contexts in which the word can occur. Another reason is based on a study in which I collected continued word associations to concrete and abstract words (de Groot, 1989). The words to be associated to were presented out of context, so I assumed the response words only to reflect context-independent information in the corresponding concepts. More responses were produced to concrete words than to abstract words, which I took to indicate that concepts corresponding to concrete words contain more context-independent information than those of abstract words do. Of course, the fewer context-independent properties in the representations, the fewer there are to be shared interlingually. This could explain the concreteness effects in bilingual processing tasks discussed in the preceding sections.

Conclusion

In the preceding sections a unitary account of performance in a number of bilingual processing tasks was suggested. Its starting point was a very simple one, namely, that conceptual representations consist of a set of meaning elements of which larger or smaller numbers may be shared by a word and its translation in another language. But it may be that things will eventually turn out to be far more complex than suggested here. Although many findings have been discussed, many others have been ignored or hardly attended to, for instance, those from studies investigating bilingual memory with episodic memory tasks. It remains to be seen whether the present framework could also account for those. Also, the present view assumes processes the details of which appear somewhat mysterious for the time being. How, for instance, are areas of (intersecting) activation in memory detected and, if necessary, evaluated? In spite of these questions, I hope to have succeeded in convincing the reader that the present view is a plausible alternative to the more established conceptions of bilingual memory and performance.

Acknowledgment
I would like to thank Lucia Dannenburg for composing the figures in this text, and Batja Mesquita for drawing my attention to a number of interesting studies on the emotion vocabulary.

References

Aitchison, J. (1987). Words in the mind: An introduction to the mental lexicon. Oxford: Basil Blackwell.

Altarriba, J. (1992). The representation of translation equivalents in bilingual memory. In R. J. Harris (Ed.), *Cognitive processing in bilinguals.* Amsterdam: Elsevier.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition, 10,* 82-93.

Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition, 11,* 211-227.

Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization.* Cambridge: Cambridge University Press.

Barsalou, L. W., & Medin, D. L. (1986). Concepts: Static definitions or context-dependent representations? *Cahiers de Psychologie Cognitive, 6,* 187-202.

Caramazza, A., & Brones, I. (1980). Semantic classification by bilinguals. *Canadian Journal of Psychology, 34,* 77 - 81.

Chen, H.-C., & Leung, Y.-S. (1989). Patterns of lexical processing in a nonnative language. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 316-325.

Chen, H.-C., & Ng, M.-L. (1989). Semantic facilitation and translation priming effects in Chinese-English bilinguals. *Memory & Cognition, 17,* 454-462.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82,* 407-428.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8,* 240-247.

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

Ervin, S. M., & Osgood, C. E. (1954). Second language learning and bilingualism. *Journal of Abnormal Social Psychology, Supplement,* 139-146.

Fischler, I., & Goodman, G. O. (1978). Latency of associative activation in memory. *Journal of Experimental Psychology: Human Perception and Performance, 4,* 455-470.

Gerard, L. D., & Scarborough, D. L. (1989). Language-specific lexical access of homographs by bilinguals. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 305-315.

Groot, A. M. B. de (1989). Representational aspects of word imageability and word frequency as assessed through word association. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 824-845.

Groot, A. M. B. de (in press). Determinants of word translation. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

Groot, A. M. B. de, & Nas, G. L. J. (1991). Lexical representation of cognates and noncognates in compound bilinguals. *Journal of Memory and Language, 30,* 90-123.

Groot, A. M. B. de, Thomassen, A. J. W. M., & Hudson, P. T. W. (1982). Associative facilitation of word recognition as measured from a neutral prime. *Memory & Cognition, 10,* 358-370.

Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart, J. L. McClelland, and the PDP Research Group, *Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 1: Foundations.* Cambridge: MIT Press.

Jin, Y.-S., & Fischler, I. (1987). *Effects of concreteness on cross-language priming of lexical decision.* Paper presented at the Southeastern Psychological Association Meeting, Atlanta, Georgia.

Katz, J. J. (1972). *Semantic memory.* New York: Harper & Row.

Keatley, C. W. (1992). History of bilingualism research in cognitive psychology. In R. J. Harris (Ed.), *Cognitive processing in bilinguals.* Amsterdam: Elsevier.

Kerkman, J. P. M. (1984). Woordherkenning in twee talen (word recognition in two languages). In A. J. W. M. Thomassen, L. G. M. Noordman, & P. A. T. M. Eling (Eds.), *Het leesproces* (The Reading Process). Lisse, The Netherlands: Swets & Zeitlinger.

Kirsner, K., Brown, H. L., Abrol, S., Chadha, N. N., & Sharma, N. K. (1980). Bilingualism and lexical representation. *The Quarterly Journal of Experimental Psychology, 4,* 585-594.

Kirsner, K., Smith, M. C., Lockhart, R. S., King, M. L., & Jain, M. (1984). The bilingual lexicon: Language-specific units in an integrated network. *Journal of Verbal Learning and Verbal Behavior, 23,* 519-539.

Kolers, P. A. (1963). Interlingual word associations. *Journal of Verbal Learning and Verbal Behavior, 2,* 291-300.

Kroll, J. F., & Curley, J. (1988). Lexical memory in novice bilinguals: The role of concepts in retrieving second language words. In M. Gruneberg, P. Morris, & R. Sykes (Eds.), *Practical aspects of memory, Vol. 2.* London: John Wiley & Sons.

Kroll, J. F., & Stewart, E. (1990). *Concept mediation in bilingual translation.* Paper presented at the 31st Annual Meeting of the Psychonomic Society, New Orleans, November, 1990.

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

Masson, M. E. J. (1991). A distributed memory model of context effects in word identification. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition.* Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Meyer, D. E., & Ruddy, M. G. (1974). *Bilingual word-recognition: Organization and retrieval of alternative lexical codes.* Paper presented at the Eastern Psychological Association meeting, Philadelphia, Pennsylvania, April.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90,* 227-234.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition.* Hillsdale NJ: Lawrence Erlbaum.

Potter, M. C., So, K-F., Von Eckardt, B., & Feldman, L. B. (1984). Lexical and conceptual representation in beginning and proficient bilinguals. *Journal of Verbal Learning and Verbal Behavior, 23,* 23-38.

Quillian, M. R. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing.* Cambridge, Mass: MIT press.

Rosch, E. (1973). On the internal structure of perceptual and semantic categories. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language.* New York: Academic Press.

Roth, E. M., & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology, 15,* 346-378.

Rumelhart, D. E., & Norman, D. A. (1985). Representation of knowledge. In A. M. Aitkenhead & J. M. Slack (Eds.), *Issues in cognitive modeling.* London-Hillsdale, NJ: Lawrence Erlbaum.

Russell, J. A. (in press). Culture and the categorization of emotions. *Psychological Bulletin.*

Scarborough, D. L., Gerard, L., & Cortese, C. (1984). Independence of lexical access in bilingual word recognition. *Journal of Verbal Learning and Verbal Behavior, 23,* 84-99.

Schwanenflugel, P. J., Harnishfeger, K. K., & Stowe, R. W. (1988). Context availability and lexical decisions for abstract and concrete words. *Journal of Memory and Language, 27,* 499-520.

Schwanenflugel, P. J., & Rey, M. (1986). Interlingual semantic facilitation: Evidence for a common representational system in the bilingual lexicon. *Journal of Memory and Language, 25,* 605-618.

Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review, 81,* 214-241.

Taylor, I. (1971). How are words from two languages organized in bilinguals' memory? *Canadian Journal of Psychology, 25,* 228-240.

Taylor, I. (1976). Similarity between French and English words—A factor to be considered in bilingual language behavior? *Journal of Psycholinguistic Research, 5,* 85-94.

Taylor, I., & Taylor. M. M. (1990). *Psycholinguistics: Learning and using language.* Englewood Cliffs, N. J: Prentice-Hall.

Tzelgov, J., & Henik, A. (1989). *The insensitivity of the semantic relatedness effect to surface differences and its implications.* Paper presented at the First European Congress of Psychology, Amsterdam, 2-7 July.
Wittgenstein, L. (1953). *Philosophical investigations.* New York: MacMillan.

Orthography, Phonology, Morphology, and Meaning
R. Frost and L. Katz (Editors)
1992 Elsevier Science Publishers B.V.

CHAPTER 21

Memory-addressing Mechanisms and Lexical Access

Kenneth Forster
Department of Psychology
University of Arizona, Tucson

The study of lexical access is important for two reasons. Not only does it deal with an integral component of the perception of language, but more generally, it raises a fundamental question about the functioning of the brain: how is previously stored information about an input pattern retrieved? The lexical domain is well suited to an experimental study of this problem, since words form a well-structured and easily manipulated set of patterns. Furthermore, due to the pioneering work of Herbert Rubenstein and his colleagues in developing the lexical decision task (Rubenstein, Garfield, & Millikan, 1971), we have access to a rich set of findings concerning the time it takes to recognize a word.

The central concept that integrates much of the theoretical and empirical work in this area is the concept of *content-addressable memory*. In this paper, we review the arguments for content-addressability, and consider the similarities and differences among the various models that have been proposed. In particular, we discuss the proposal that the mental lexicon is only approximately content-addressable, and that serial search mechanisms are inevitably involved in lexical access. We will discuss some of the evidence in favor of this claim, and deal with some of the objections to the notion of serial search.

Content-addressable retrieval

In a conventional computer memory, each memory location is assigned a number, which represents its address. Storage or retrieval of data from a particular memory location requires that the bit pattern corresponding to its address be first loaded into the address decoder, a circuit which selects the memory location designated by that address. Once selected (i.e., enabled), the data within this memory cell can be retrieved or modified. This retrieval function, *R*, can be expressed as follows:

$$R(address) = contents$$

For this reason, this type of memory is referred to as location-addressable memory. With content-addressable memory, however, what is retrieved is the address of the memory cell that has a specified content. That is,

$$R(contents) = address$$