

# Interpreting in Linguistics

Richard Hudson

## 1 Interpreting and translating

In linguistics, the word *interpret* is not a technical term (though its derived forms *interpreted* and *interpretive* are; their meanings are very technical and irrelevant here). Consequently, I shall be using it here in its everyday meaning. But what is its everyday meaning? As an exercise in lexical semantics, I would suggest a definition along the following lines:

(1) Interpreting X is expressing X in a way that makes it easier to understand.

I believe (1) captures most if not all of the ways in which we use this term in everyday life. For example:

- If I interpret your behaviour, I explain it (possibly in words).
- If I interpret your dreams, I explain them or look for a hidden meaning.
- If I interpret your words monolingually, I look for hidden meanings.
- If I interpret your words bilingually, I translate them into another language.

## 2 Meaning and function

Now suppose you hear me say: [sɪəʊbʊk]. (For those not familiar with the International Phonetic Alphabet, [ə] is the vowel in *book*, [ɹ] is the English ‘r’ as in *red* or *very* – a continuous sound which easily turns into a [w], in sharp contrast to the various ‘r’ sounds found in languages such as French, German, Italian and Spanish – and [ɒ] is the vowel in *top*. All the other letters correspond to the sounds you would expect.) How can you interpret this? In other words, what does interpreting mean in this kind of situation? You might push for a deeper understanding by asking yourself why I said it. In some rather unusual situations, I might have said it simply to test your grasp of the IPA; linguistics students generally undergo ‘ear training’ in which the aim is to transcribe ‘nonsense words’ which are uttered purely for the sake of their sounds. This is why ear trainers say words, and students know it; but this is a very untypical situation. In everyday life, if I say something to you it’s in order to affect your thinking – in short, to make you think some particular thought. Let’s call this thought X. The challenge for everyday interpreting of everyday language is to work out what X is.

The question, therefore, is what is X when I say [sɪəʊbʊk]? What thought does English link [sɪəʊbʊk] to? This is a complicated way of asking what [sɪəʊbʊk] **means**, or in slightly more technical terms, what its **function** is – what it is used for. *Function* is actually a more general term than *meaning*, at least in the latter’s everyday sense, because we tend to assume that meaning must have something to do with the world, whether we call this ‘reference’ or ‘denotation’ or ‘description’. If we ask what is the meaning of the word *book*, we are looking for a description of the kind of thing this word can refer to; so

we would not be satisfied with information about the word's grammar: that it is a singular common noun. But words are in fact linked to grammatical information of this kind, and when we hear them they do indeed make this information active in our minds, so a thorough interpretation of a word should take a broader perspective. The notion we need, therefore, is not 'meaning' but 'function'.

Units of language may have all sorts of different functions. Many of them are covered by the following list, which starts with the familiar notion of meaning and moves towards increasingly concrete units. In each case, a more concrete unit has the function of helping to manifest more abstract units.

- **Lexical-semantic** functions

This is the everyday notion of 'meaning'; e.g. the function of the word *seal* is to mean 'Seal<sub>animal</sub>' or 'Seal<sub>stamp</sub>' – i.e. it refers either to a kind of animal, or to a thing like a stamp.

- **Inflectional** functions

The function of a morphological inflection is to modify word meaning (and word syntax). E.g. in the inflected word *seals*, the suffix {s} signals 'plural'.

- **Morphological** functions

Each part of a word has some function inside that word which defines the contribution that it makes to the word's grammar and meaning. Compound words provide a simple example of this, because they usually have two parts, of which one is the 'head', which defines the basic meaning, and the other is the 'modifier', which modifies this to make it more specific; e.g. the morpheme {oil} is the head in the word *seal-oil*, which refers to a kind of oil made from seals.

- **Morpho-phonological** functions

Morphemes are 'realised' (in the French sense of 'made real') by sounds, so conversely the function of a sound is to help realise morphemes – to make those morphemes audible, and to distinguish them from other morphemes. For example, the sounds /s/, /z/ or /ɪz/ realise the morpheme {s} (as in *cats*, *dogs* and *horses* respectively).

- **Phonological** functions

However, sounds have other kinds of function as well, one of which is to help build syllables and other phonological structures. Syllables are often divided into an onset, a nucleus and a coda, illustrated respectively by the three sounds in *book*. Thus the /b/ sound in *book* has not only the morpho-phonological function of helping to realise the morpheme {book}, but also the phonological function of providing the syllable's onset. Similarly in our example [sɪəʊbʊk], [ɔ] is the coda of the first syllable (an odd status, in fact, since it is a vowel and codas are usually consonants). This introduces another kind of phonological function: sounds can be classified in terms of quite abstract categories such as 'vowel' and 'consonant', which determine the way they are used, so being a vowel or a consonant is another kind of function. A functional classification can be quite different from a purely phonetic one. For example, the sound [ɔ] is functionally a vowel in *book*, but in modern English (especially in London) it is functionally a consonant in a word such as *feel* or *well*, where [ɫ] (the English 'dark l') has recently 'vocalised' into [ɔ].

- **Phono-articulatory** functions

We make sounds by putting our vocal apparatus – our lips, tongues and so on – into particular physical configurations, and blowing. This means that each sound is realised by some configuration of the vocal tract; for example, to make [ɔ] we round our lips, raise the back of our tongues and let our larynx vibrate. Each of these arrangements has the function of helping to make this sound.

The above list showed how the familiar notions of words and word meanings were part of a chain of functional links which ultimately involve physical actions. If we had been considering written language we would have followed a somewhat different path which would have ended in physical marks on the page. At the other end of the chain, however, we can go beyond the lexical-semantic functions with which we started. In increasing order of abstractness, we can distinguish the following functions:

- **Syntactic functions**

Every word in a sentence has a syntactic (or grammatical) function. This is a very traditional notion, and we have inherited a useful set of grammatical functions such as ‘subject’ and ‘object’ which date ultimately back to the Greek grammarians. For example, in (2) *he* is the subject and *the seal* is the object.

(2) He saw the seal.

This information plays an important part in understanding the sentence, because it tells us who did what. From what we know about the meaning of the verb *see*, we know that seeing involves a pair of eyes (a ‘see-er’) and a visible entity (a ‘see-en’); and from its grammar we know that the subject is the see-er and the object the see-en. (More accurately, it is the people or things referred to by the subject and object that have these semantic roles.) Thus a word’s grammatical function indicates, among other things, what semantic role it plays.

- **Semantic functions**

Semantic functions are the roles that meanings play in the total meaning of a sentence, and as explained in the previous paragraph, they can generally be ‘read off’ the syntactic functions of the words concerned. In the case of verbs and their dependents, analysts often use categories such as ‘agent’, ‘patient’, ‘experiencer’ and ‘stimulus’, and in those terms the person mentioned in (2) is the experiencer while the seal is the stimulus. However it is not just verbs that have semantic functions: virtually every word does, because each word contributes some bit of meaning which has to be related to the meanings of other words: for example, in *John’s hat*, John is owner of the hat, in *behind John*, he is the ‘landmark’ on which the position is based, in *the book*, the article shows that the book is already known, and so on and on. These relations in semantic structure are semantic functions.

- **Pragmatic functions**

Finally we have the function of the meaning, which is to influence the hearer’s mind in a particular way. If we want to make the hearer think of John, we have a wide range of linguistic options including *John* and *he*, each of which has a different meaning; different meanings have different side-effects on the hearer, but they all have the same main effect of making the hearer think of John. Conversely, a single linguistic form may refer to different people or things on different occasions; e.g. *me* refers to whoever happens to be speaking, and *the seal* refers to whichever particular seal is under consideration at the

moment. Similarly, we can use a single linguistic form to achieve many different effects; for example, consider (3):

(3) The door's open.

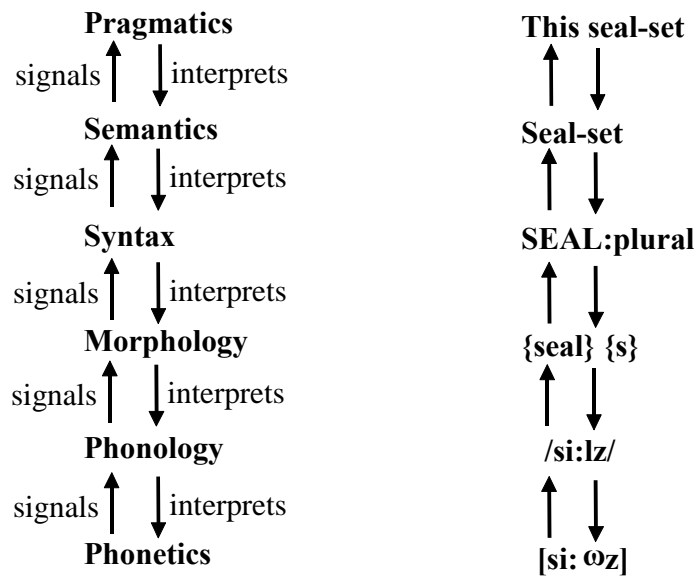
This might be a description (What's the house like?), an invitation (Can I come in?) or a request for action (Can I do anything for you?).

### 3 Signalling and interpreting

To summarise the discussion in section 2, language is traditionally (and correctly) seen as organised in 'levels' (or 'levels of analysis'). Different linguists have different views on what the levels are, but the following list is fairly uncontroversial and is the one I have espoused in other work (Hudson 1990; see also the Encyclopedia of English Grammar and Word Grammar at [www.phon.ucl.ac.uk/home/dick/enc-gen.htm](http://www.phon.ucl.ac.uk/home/dick/enc-gen.htm)):

- Pragmatics
- Semantics
- Syntax
- Morphology
- Phonology
- Phonetics

The levels are organised hierarchically, and conventionally meaning is at the top and phonetics at the bottom. In the hierarchy, each level 'interprets' the one below it in the sense that I give to *interpret* in (1). Thus the pragmatic analysis makes the semantic one easier to understand because it answers the question: why did he say something with that meaning? The semantic analysis makes the syntactic one easier to understand by explaining why the words were selected and organised as they were; and so on down the list. The converse of the 'interpret' relation is the one I called '**realise**' above, but which could also be called '**signal**': semantic structure signals pragmatic structure, syntactic structure signals semantic structure, and so on. This chain of interpret/signal relations is shown in Figure 1.



**Figure 1** *The hierarchy of language levels*

The column on the left shows the chain of relations in the abstract, while the one on the right gives a concrete application. It invites you to imagine a scene in which someone says the sounds [siɔz] (where you will recognise the first three sounds of our target [siɔɪk]) and achieves the pragmatic effect of referring to a particular colony of seals (e.g. those seen on a trip to the zoo). The intervening chain explains how these particular sounds achieved that particular pragmatic effect.

The main point of the example, as of the preceding discussion, is summarised in (4):

- (4) Interpreting language is a multi-stage process in which relatively concrete elements are interpreted in terms of relatively abstract ones.

There is no single structure which could be called ‘the interpretation of the sentence’, nor is there a single structure which is interpreted.

#### **4 Decoding or interpretation?**

How does interpretation of language work? One possibility is that it works quite differently from level to level, and one popular view is that ‘language proper’ works in one way, and the last step (from semantic to pragmatic structure) is different (Sperber, Wilson 1995: 9). According to this view, language proper is a ‘code’, just like morse code, in which interpretations can be assigned mechanically, whereas a pragmatic interpretation requires intelligent use of contextual information. This view strikes me as wrong, not because I think pragmatics is also code-like, but because I think the whole of language requires intelligent use of contextual information. My evidence is that at all levels, interpretation confronts ambiguities and uncertainties, and at all levels contextual information plays an important part in resolving these uncertainties. We may not be

aware of the intelligent work involved (say) in mapping a string of sounds onto a particular word, but that's just because we generally do it so well; but we do notice failures which show how far from simple coding the process is.

In short, language cannot be a code because if it was, communication would take place much more smoothly than it in fact does. Moreover, it would be easy to program a computer to interpret speech automatically; but in fact this is extremely hard. Although automatic speech recognition now works on our telephones for very limited tasks, the technology still lags way behind human achievement, and the same is true even for systems that use computers to interpret written language. None of these difficulties would arise if language was merely a code.

The uncertainty that arises in interpreting language seems to have three sources: ambiguity, variation and linking.

- **Ambiguity**

This is caused by one form having several functions, and is by no means confined to the 'top' levels in the hierarchy. At last I can reveal the point of my initial example of the sounds [siɔɹɒk]: these sounds can in fact be interpreted in several different ways, as shown in (5) through (8):

- (5) This [siɔɹɒk] (= seal rock) is home to a lot of seals.
- (6) This [siɔɹɒk] (= sea or rock) in your picture is blurred – I don't know if it's sea or rock.
- (7) This [siɔɹɒk] (= sea'll rock) the boat.
- (8) This [siɔɹɒk] (= seal rock) is a very good rock for making seals out of.

Most of us are not sufficiently aware of phonetics to notice that the same pronunciation fits *seal rock*, *sea or rock* or *sea'll rock*; nor are most of us likely to imagine that the compound *seal rock* could have two different meanings. There is a vast experimental literature showing that the context pushes us towards the correct interpretation with little or no conscious thought (Harley 1995: 158), so it is not surprising that we think it happens automatically.

- **Variation**

This is the converse of ambiguity, in which several alternative forms share the same function. The alternatives may correlate with different (or no) contextual influences:

- Variation with speaker (accent/dialect), e.g. *cup* as said in Sheffield or London, or Yorkshire *tret* contrasting with southern and standard *treated*.
- Variation with situation (formality etc), e.g. *try* or *attempt*.
- Free variation, e.g. *someone* or *somebody*.

- **Linking**

This involves the way in which one element is linked to other words or things, and can be found at many different levels, for example:

- Syntax: every word is linked to other words, so the hearer or reader has to decide which words in a sentence are linked to which, e.g. does (9) mean that he talked in English or about interpretation in English?

(9) He talked about interpretation in English

- Pragmatics: how are words linked to things (or people), e.g. who does *he* in (9) refer to?

These three kinds of complication or uncertainty combine freely with each other at every level of language, and explain not only why automatic processing of language by computer is such a gigantic task, but also why even monolingual communication is so fraught with danger and difficulty, to say nothing of the even greater dangers in communicating across languages.

## 5 How do we interpret?

Given these enormous problems, it is impressive that normal everyday communication proceeds so painlessly most of the time. How do we do it? What follows is the conclusion that I have arrived at after several decades of intense work on language structure, some reading in psycholinguistics, some practical experience with programming and a great deal of thought. It is part of a general theory which I have called ‘word grammar’, which is highly controversial but not entirely out on a limb – it fits very comfortably into a recent trend in linguistics called ‘cognitive linguistics’ (Barlow, Kemmer 2000; Croft, Cruse 2004; Tomasello 1998), and most of its parts are more or less accepted in cognitive psychology and/or artificial intelligence. (For more information about it, including published references, see my web site at [www.phon.ucl.ac.uk/home/dick/wg.htm](http://www.phon.ucl.ac.uk/home/dick/wg.htm).)

I believe that all the elements of language are represented by nodes which are connected to each other in a giant network; moreover, this network of language elements is itself contained in an even vaster network which contains all our (propositional) knowledge about the world, including transient contextual knowledge as well as permanent knowledge. In this theory, each piece of experience has a temporary node in the network, and interpreting the experience involves finding models for it among the permanent nodes. Let’s call the experience X, and suppose it is the experience of hearing the sounds [siəz] in the context of a visit to the seal enclosure in a zoo.

Now we come to the crucial fact about cognitive networks: they are highly active, thanks to electrical energy which activates individual nodes, and which spills over from highly active nodes to their neighbours – a very crude process which drives our thinking processes and which is called **spreading activation**. Activation is directed by attention, so we direct it towards the nodes for people and things in which we are interested. In the present example, we are interested in:

- X, the sounds that we have just heard
- the meaning which we know these sounds must have, although we don’t at first know what it is; let us call it Q (for ‘query’)
- the seals that we are looking at.

Consequently, the nodes for X, Q and the seals will be highly active, and by spreading activation their activity will spill over onto neighbouring nodes. However, the energy supply is limited so energy dies down quickly – except on nodes which receive energy from more than one source. This selective activation of nodes is the key to understanding. In our example, a few nodes will receive activation from both X, Q and the seals, so they

will remain active and define collectively the most active **global interpretation**, which combines activity from both X and Q; and thanks to two extra processes called binding and inheritance the target meaning is enriched by inheriting from the most active suitable node, the one for ‘seal’. This process is illustrated schematically in Figure 2, where the nodes are shown as boxes whose border thickness shows the level of activation. The diagonal arrows show alternative interpretations which could have been chosen, but the correct interpretation is supported from both ends of the hierarchy; for example, the meaning ‘seal’ (in the top left corner) is heavily active thanks to its close link to Sammy the Seal, but it is also supported by activation via the word *seal* whose meaning is a set of seals.

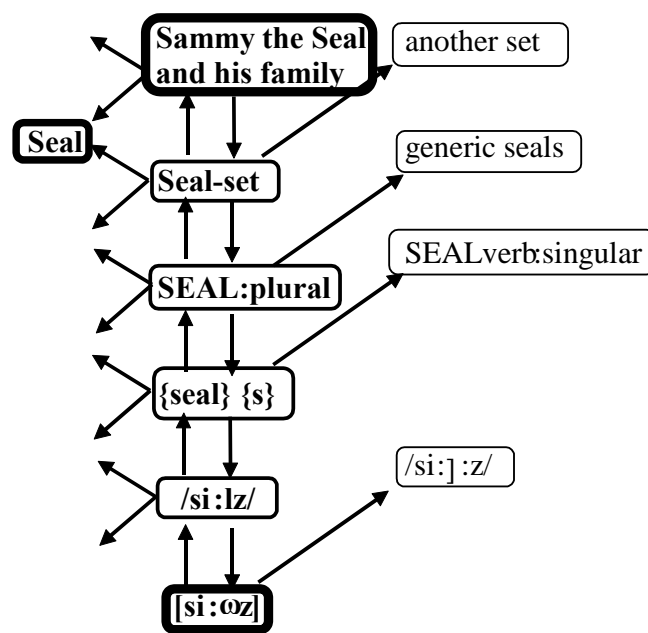


Figure 2 *Interpreting sounds in terms of meanings*

The crucial part of this account is the rather mundane observation that understanding a piece of experience is not a one-way process, but is heavily influenced by the context. In this case the context is of course the intellectual context of concepts which are currently active, but since these are part of the same mega-network as the language, their activation can help to select a sensible interpretation for the experience.

Rather conveniently, the same theory applies with very little change to speaking, in which we have to select behaviour (speech activity or writing) which fits the message we want to express. Fortunately we can use the same linguistic network in both directions, and we use it in much the same way, with some activation coming from both ends of the linguistic hierarchy: from our intended message, but also from the target word. This must already exist in our minds because it defines our aim: to say or write a word that has the intended meaning. We start with a schematic representation of this



word which, being our current goal, is highly active; and activation spreads out as before from both of these poles, defining a global path between them. The procedure for selecting the sounds [siɔz] in the zoo context is illustrated in Figure 3, showing activation once again flowing from both Sammy the Seal and the target sounds. (Like the previous diagram, this one ignores a more complex process in which the target starts as a schematic node which is gradually enriched by inheriting properties from other nodes.)

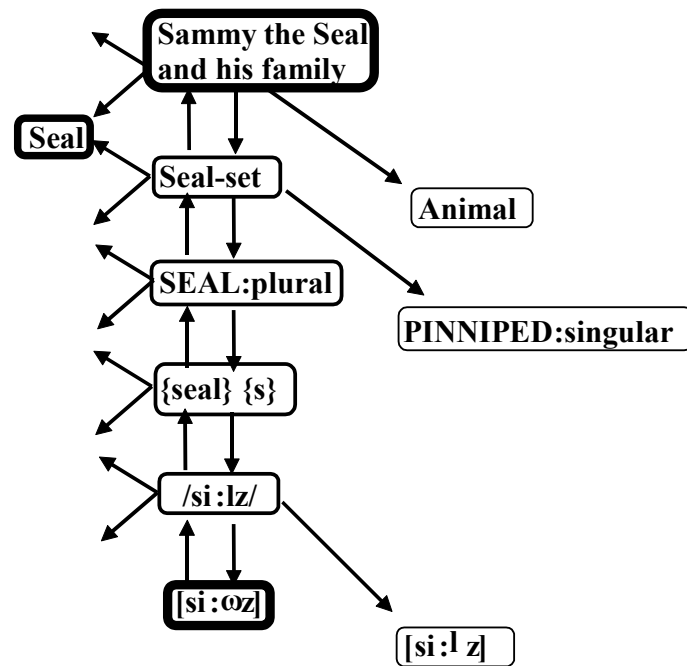


Figure 3 *Expressing meaning as sounds*

## 6 Interpreting as translation

I believe these general ideas about understanding and speaking throw some light on the process of translating from one language to another, which has a special interest of its own in that it combines interpreting with speaking: the translator interprets the source language text and speaks or writes its equivalent in the target language. Even more interestingly for linguists and psycholinguists, it involves two different languages so the highly automatic behaviour of a skilled conference interpreter will be hard to explain if different languages are considered to live in different modules of the mind, as some people have suggested (Romaine 1989:84).

In a network model, both the languages are part of the same network so it is very easy for them to ‘communicate’ with each other. If we interpret communication in terms of spreading activation, this means that an active node in one language may easily activate a related node in the other; for example, when a French-English bilingual hears the word *seal*, the node for its French equivalent (*phoque*) may be slightly activated even in monolingual conversation, but when the aim is to translate into French, it may be very heavily activated. This is illustrated in Figure 4, where translation is shown going ‘the

long way round' via the semantics, albeit not the even longer way round via pragmatic interpretation. It is very easy to imagine connections forming between *seal* and *phoque*, after years of translation practice, at a lower level – maybe even at the phonological level. The more such links exist, the more activation will spread from the English word directly to the French one and the faster (and easier) will be the translation process.

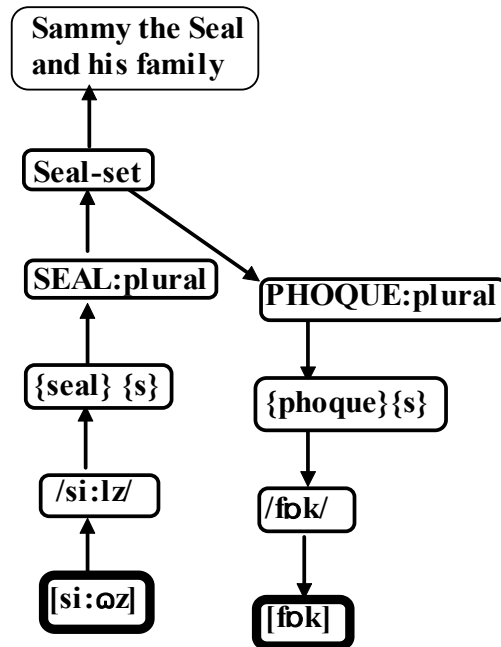


Figure 4 Translating from English to French

## 7 Some questions

The above remarks were based on my experience of thinking about language in general, rather than from any kind of expertise in theorising about translation as such. However, as I suggested above, this process clearly holds some interest for linguists and psycholinguists, so I should like to end with some questions which seem to me to deserve research (and which, for all I know, may already have received this attention). They may be old hat to those in the field of translation and interpreting, but even old questions often turn out to hold new interest in the light of new theories.

- How much does the translator need to know about the pragmatics of what is being translated? Clearly some things are irrelevant – e.g. *seal* will probably be translated as *phoque* regardless of which particular seal is being referred to. On the other hand, context is important to the extent that it removes ambiguities. Are there any general principles for distinguishing the relevant from the irrelevant?
- How much does the activation of one language spread to the other? This picks up the suggestion made above that equivalent items in the two languages may be linked directly, rather than only via their shared meaning, and that any such links will play an important role in translation by allowing activation to spread directly.

- Is this activation greater in a skilled simultaneous interpreter? If direct links exist at all, we would expect them to be based on experience of translation so the skill of an expert interpreter will be reflected in a particularly rich collection of cross-language links. Psychologists have standard experimental paradigms for measuring activation (Harley 1995:17), so this question could in principle be answered by experimental work.
- Finally, what does ‘accurate translation’ mean? No doubt this question is one that translators and interpreters visit often, but the network view may suggest a new range of possible answers in terms of effects on the general cognitive network. One aspect of the network view which I have not mentioned here is that mental networks (obviously) belong to individuals and are based on these individuals’ experiences, so there is no reason at all to expect different individuals to have the same mental networks, however similar their experiences and culture may be. This means that communication within one language is never perfect – my concept for ‘seal’ is probably slightly different from yours, so our minds will never meet perfectly. This monolingual pessimism offers a useful perspective on cross-lingual communication, because the question then is not whether a translation can be ‘accurate’, but whether the communication it achieves is any less successful than the partial success that we expect in monolingual communication.

Richard Hudson <dick@linguistics.ucl.ac.uk>

## References

- Barlow, Michael and Suzanne Kemmer. 2000. *Usage Based Models of Language*. Stanford: CSLI.
- Croft, William and Alan Cruse. 2004. *Cognitive Linguistics*. Cambridge: Cambridge University Press.
- Harley, Trevor. 1995. *The Psychology of Language*. Hove: Psychology Press.
- Hudson, Richard. 1990. *English Word Grammar*. Oxford: Blackwell.
- Romaine, Suzanne. 1989. *Bilingualism*. Oxford: Blackwell.
- Sperber, Dan and Deirdre Wilson. 1995. *Relevance. Communication and cognition*. Oxford: Blackwell.
- Tomasello, Michael. 1998. *The New Psychology of Language. Cognitive and functional approaches to language structure*. Hillsdale, NJ: Erlbaum.