

Semantic representation in LCCM Theory

Vyvyan Evans

[To appear in *New Directions in Cognitive Linguistics*, ed. By V. Evans & S. Pourcel. John Benjamins.]

1. Introduction

In this paper I am concerned with the approach to semantic representation adopted in the *Theory of Lexical Concepts and Cognitive Models*, or *LCCM Theory* for short.¹ LCCM Theory takes its name from the two central theoretical constructs adopted in the theory: the *lexical concept* and the *cognitive model*. A lexical concept is a component of linguistic knowledge, the semantic pole of a *symbolic unit* (in Langacker's e.g. 1987 terms), and encodes a bundle of different types of linguistic knowledge. Put another way, the lexical concept represents the means adopted in LCCM Theory of modelling units of *semantic structure*. In contrast, a cognitive model is a component of conceptual knowledge, which is to say, non-linguistic knowledge. Hence, the cognitive model represents the means adopted in LCCM Theory of modelling units of *conceptual structure*. LCCM Theory assumes that lexical concepts and cognitive models are types of semantic representation belonging to two distinct representational systems, which have distinct and divergent functions. These are the *linguistic system*, which encodes semantic structure and the *conceptual system* which encodes conceptual structure. Following arguments presented by Barsalou et al. (To appear), I suggest that the linguistic system evolved, in part, by facilitating more effective control of the extant representations in the conceptual

¹ LCCM Theory represents an attempt to provide a cognitively realistic account of the semantic mechanisms involved in the construction of meaning in language understanding.

system. That is, linguistic representations are specialised for providing a ‘scaffolding’ to structure conceptual representations, thereby facilitating their use in communication (cf. Talmy 2000). While the much older — in evolutionary terms — conceptual system evolved for action and perception, i.e. for non-linguistic purposes, the emergence of language facilitated the use of conceptual representations in linguistically-mediated meaning construction, thereby providing modern humans with a significant evolutionary advantage. With the association of linguistic and conceptual representations, humans were able to engage in the advanced symbolic behaviours that led to the explosion of sophisticated ritual practice, material culture, art and science around 50,000 years ago during the period known as the Upper Palaeolithic (Mithen 1996).

My argument, in a nutshell, is this: the semantic representations in the linguistic and conceptual systems interact for purposes of linguistically mediated communication. Together, the lexical concept and the cognitive model form a level of representation that I refer to as *semantic representation*. My purpose in this paper is to describe the nature of the lexical concept, the nature of the cognitive model, and the nature of the interaction between the two.

In order to better illustrate these distinctions, and how they intersect, figure 1 provides a diagrammatic representation of the relationship between the linguistic and conceptual systems, as assumed by LCCM Theory. Figure 2 provides a illustration of semantic representation in LCCM Theory. In figure 1 the dashed line between the lexical concept in the linguistic system and the cognitive model (represented by the circle) in the conceptual system represents a path of access relating the two. Figure 2 is the same as figure 1 except that it additionally features a dashed ellipse encircling the lexical concept (in the linguistic system) and the cognitive model (in the

conceptual system), the two types of representations which collectively comprise semantic representation.

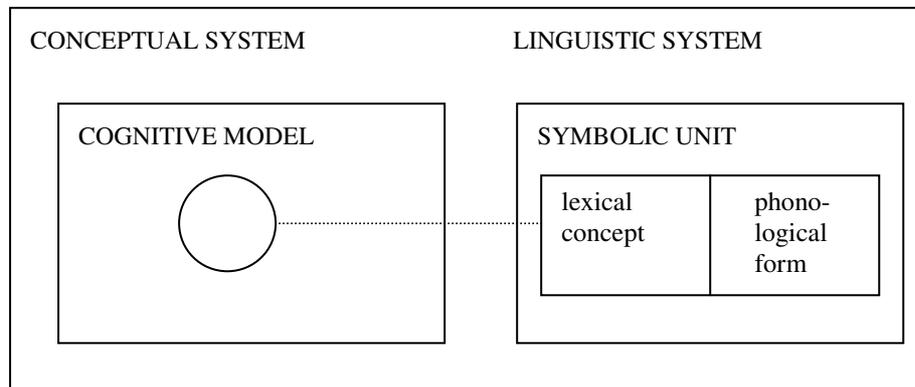


Figure 1. The relationship between the linguistic and conceptual systems

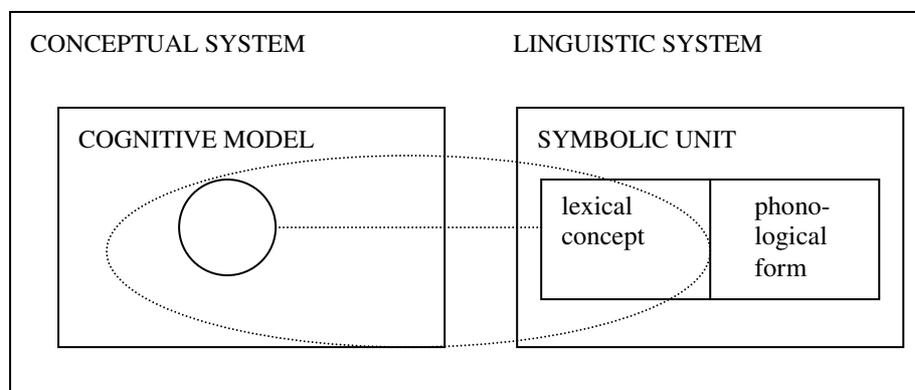


Figure 2. Semantic representation in LCCM Theory

2. The distinction between linguistic content and conceptual content

According to Talmy (2000), a central design feature of language is that the concepts expressed are divided into two subsystems. Talmy characterises this in terms of what

he refers to as the *grammatical subsystem* and *lexical subsystem*. These two subsystems serve to express what we might refer to as the experiential complex — Talmy uses the term: *cognitive representation* — that a speaker attempts to evoke in the listener by virtue of deploying language. The range of concepts expressed by the grammatical subsystem is highly restricted cross-linguistically, providing a basic framework for the structuring of the experiential complex that language users seek to evoke in their interlocutors. Put another way, the lexical concepts associated with the grammatical subsystem have *schematic content*, providing a structuring function. Thus, lexical concepts of this sort provide a ‘scaffolding’ so to speak, across which the rich content associated with lexical concepts of the lexical subsystem can be draped. In contradistinction to this, the lexical concepts associated with the so-called lexical subsystem provide *rich content*, giving rise to the details (rather than structural aspects) of the cognitive representation.

An important aspect of Talmy’s work is the claim that the distinction between rich versus schematic content corresponds to a bifurcation between form types: open-class versus closed-class forms. Closed-class forms are so-called because it is considered more difficult to add members to this set. This set of lexical items includes the so-called ‘grammatical’ or ‘function’ words such as conjunctions, determiners, pronouns, prepositions, and so on. In contrast, open-class forms include words belonging to the lexical classes: noun, verb, adjective and adverb.

While the concepts expressed by closed-class forms encode schematic content, they are nevertheless essential for the expression of the cognitive representation. To make this point clear, consider the following semantic analysis of the range of open- and closed-class elements which comprise the utterance in (1):

(1) **A waiter served the customers**

The forms in bold: **a**, **-ed**, **the** and **-s** are associated with the grammatical subsystem. Their semantic contribution relates to whether the participants (waiter/customers) in the experiential complex evoked by (1) can be easily identified by the hearer (the use of the indefinite article **a** versus the definite article **the**), that the event took place before now (the use of the past tense marker **-ed**), and how many participants were involved (the absence or presence of the plural marker **-s**).

In contrast, the forms in italics: *waiter*, *serve* and *customer* are associated with the lexical subsystem. That is, their semantic contribution relates to the nature of participants involved in the experiential complex, and the relationship holding between them. In other words, while the closed-class forms encode content relating to structural aspects of the experiential complex evoked, the open-class forms are associated with detailed information concerning the nature of the participants, scenes involving the participants, and the states and relationships that hold.²

I argue that the distinction in content evoked by language, and pointed to by Talmy, relates to a distinction in content associated with the linguistic system (and lexical concepts) on the one hand and the conceptual system (and cognitive models) on the other. The two distinct types of content implicated I refer to as *linguistic content* and *conceptual content*.

Dealing with the latter first, conceptual content relates to the rich content evoked by open-class forms. Information of this kind is primarily perceptual in nature.

² The closed-class forms mentioned thus far all have an overt phonetic realisation. However, each of the examples discussed also include closed-class forms that are phonetically implicit. Examples of phonetically implicit forms include lexical classes: e.g. noun, verb; lexical subclasses: e.g. count noun, count noun; grammatical relations: e.g. subject, object; declarative versus integrative forms, active voice versus, passive voice, and clause-level symbolic units such as the ditransitive construction, and so forth.

By ‘perceptual’ I have in mind information that derives from i) *sensory-motor systems*—those sensory systems that recruit information relating to the external environment and the human individuals’ interaction with the environment — as well as ii) *proprioception* — the systems that recruit information relating to the motor aspects of the body’s own functioning — and iii) *subjective experience* — which includes experiences ranging from emotions, temporal and other cognitive states, to the visceral sense, all discussed in more detail later. Accordingly, I am following Barsalou (1999) in defining perceptual experience more broadly than has traditionally been the case.

Conceptual content provides records of perceptual states, in the sense just given. As such, it is *analogue* in character: it re-presents perceptual information that parallels the multimodal perceptual experience that it constitutes a representation of.³ As such, conceptual structure is not suitable for being encoded in language. After all, language as a representational system, consisting of symbolic units, is simply not equipped to directly encode the rich, multimodal aspects of perceptual experience.

In contrast, I argue that the schematic content identified by Talmy is not an analogue representation of perceptual experience. Rather, it represents an abstraction over perceptual content of various sorts, provided in a form that can be encoded in language, i.e. by lexical concepts. Content of this kind constitutes what I refer to as linguistic content, and forms part of the information encoded by a lexical concept.

While the distinction between rich and schematic aspects of the cognitive representation provide the basis for my distinction between linguistic and conceptual content, the distinction in open-class and closed-class forms provides evidence for a closely related distinction concerning the nature of lexical concepts — recall that

³ Conceptual content is not an exact record of the perceptual states that are captured. Rather, it is somewhat attenuated. See Barsalou (1999) for discussion.

lexical concepts are conventionally associated with phonological vehicles (i.e. forms). The distinction in vehicle types provides evidence that lexical concepts fall into two distinct categories. Closed-class vehicles are associated with lexical concepts which are specialised for encoding linguistic content. Lexical concepts of this sort I refer to as *closed-class lexical concepts*. Open-class vehicles, while also encoding linguistic content, are, in addition, specialised for serving as *access sites* to conceptual content. Lexical concepts of this sort I refer to as *open-class lexical concepts*. This distinction is captured in figure 3.

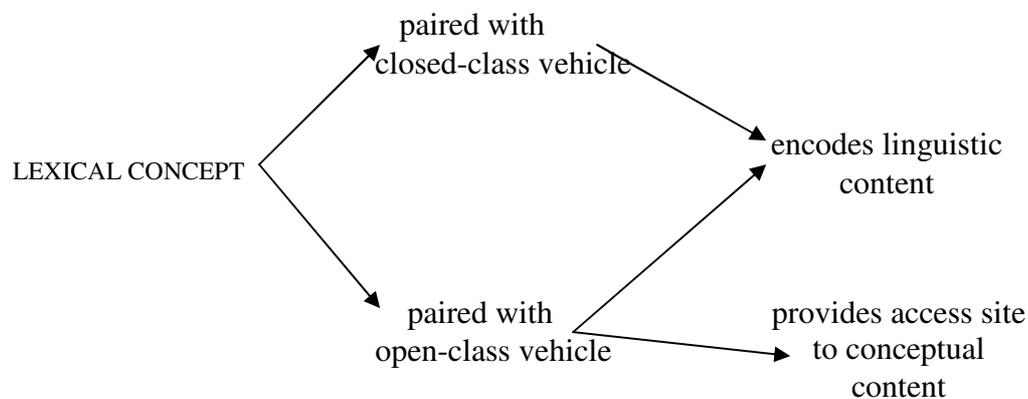


Figure 3. The distinction in content associated with lexical concepts

The distinction between the terms ‘encode’ and ‘afford access’ in the previous paragraph is critical here. Linguistic content is encoded by lexical concepts precisely because this is the content which makes up lexical concepts. However, conceptual content is associated with a different representational type, the cognitive model, which is non-linguistic in nature. Thus, conceptual content is not directly encoded in language, although the linguistic system has developed the means to access conceptual content, discussed in more detail later in the paper. Table 1 provides a

summary of the way some of the key terms introduced so far are used in LCCM Theory.

Table 1. A summary of key terms in LCCM Theory

Term	Description
Linguistic system	The collection of symbolic units comprising a language, and the various relationships holding between them
Symbolic unit	A conventional pairing of a phonological form or vehicle and a semantic element
Lexical concept	The semantic element that is paired with a phonological vehicle in a symbolic unit
Linguistic content	The type of content encoded by a lexical concept. This content is of a highly schematic type that can be directly encoded <i>in</i> language
Conceptual system	The body of non-linguistic knowledge captured from perceptual experience that is made of perceptual states. This knowledge derives from sensory-motor experience, proprioception and subjective experience
Cognitive model	The representational form that knowledge in the conceptual system takes, as modelled in LCCM Theory. Consists of frames which give rise to a potentially unlimited set of simulations
Conceptual content	The nature of the knowledge encoded by a cognitive model
Lexical representation	The primary substrate deployed in linguistically-mediated meaning construction, and modelled in terms of symbolic units and cognitive models
Semantic representation	The semantic dimension of lexical representations, consisting of semantic structure and conceptual structure
Semantic structure	That part of semantic representation encoded by the linguistic system. Semantic structure is modelled, in LCCM, Theory, by lexical concepts,
Conceptual structure	That part of the semantic representation encoded by the conceptual system. Conceptual structure is modelled, in LCCM Theory, by cognitive models

3. An illustration

Before proceeding further, I provide a brief illustration of the distinction between linguistic and conceptual content. Consider the use of the lexical item *red* in the following examples, adapted from Zwaan (2004):

- (2)
 - a. The teacher scrawled in red ink all over the assignment
 - b. The red squirrel is in danger of becoming extinct in the British isles

Zwaan makes the point that in linguistic examples such as (2), *red* designates two different sorts of sensory experience. That is, while the hue derived from the use of *red* in (2a) is quite a vivid red, the hue prompted for by (2b) is likely to be closer to a dun/browny colour. That is, what I refer to as the *semantic potential* of *red* is not ‘there’ in the word itself. That is, whatever *red* designates, we are not dealing with purely linguistic knowledge. Rather, the word *red* provides access to perceptual information and knowledge, which can be reconstructed or *simulated* — I will have more to say about this idea below. Put another way, the hue derived is not a function of linguistic knowledge, but relates to what I am referring to as conceptual content. This is not to say that *red* does not provide linguistic knowledge. My point is that the vehicle *red* has an associated lexical concept that I gloss as [RED] — in LCCM Theory lexical concepts are glossed by a term, here ‘red’, in small capitals inserted in square brackets. This encodes schematic linguistic content, designating that an entity is being referred to, that the entity being referred to is a relation of some kind, and that the relation is specifically an attribute of a thing. In short, while linguistic content includes highly schematic semantic knowledge, conceptual concept concerns richly detailed perceptual knowledge.

4. The nature of semantic structure

In LCCM Theory semantic representation consists of units of semantic structure and conceptual structure and their interaction. In this section I address, in more detail, the nature of semantic structure.

Semantic structure is modelled in terms of the theoretical construct of the lexical concept, which constitutes a unit of semantic structure. As noted above, lexical concepts encode linguistic content. Linguistic content represents the informational form that conceptual structure takes for *direct* representation *in* language. Put another way, linguistic content takes a form that can be encoded in a format that is externalised in an auditory stream (or a manual gestural stream in the case of signed language), which is time-pressured. Such a format presumably requires filtering out the complexity associated with the range of perceptual experiences — in the sense defined above — encoded. Hence, a lexical concept can be thought of as a bundle of different types of highly schematic content which is thereby specialised for being encoded in language. I detail some of the key aspects of linguistic content below.

4.1 Parameters

One way in which knowledge, in general terms, can be represented is in terms of richly inflected nuances that serve to reflect the complexity of experience. An alternative way is to ‘compress’ such fine distinctions into two, three or more, much broader, and hence, far more general distinctions. These I refer to as *parameters*. Linguistic content serves to encode content by adopting the latter strategy, which is to say, to employ *parameterisation*. Parameters are hence part of the bundle of information that a lexical concept serves to encode.

To illustrate this notion, consider the complex range of expressions that a language user might employ, in English, in order to ‘locate’ themselves with respect to time, thereby facilitating time-reference. Any one of the following could conceivably be employed, depending upon context: *today, January, 2008, the day*

after yesterday, the day before tomorrow, this moment, now, this second, this minute, this hour, today, this week, this month, this quarter, this year, this half century, this century, this period, the 8th day of the month, this era, this millennium, and so on.

In contrast, parameterisation functions by dividing all the possible permutations relating to a given category, such as time-reference, into a small set of divisions: parameters. Such parameters might distinguish between the past, for instance, and the non-past. Indeed, this is the basis for the tense system in English, as illustrated by the following:

- | | | |
|-----|-----------------------|----------|
| (3) | a. He kicked the ball | Past |
| | b. He kicks the ball | Non-past |

English encodes just two parameters that relate to Time-reference: Past versus Non-past, as exhibited by the examples in (3), and thus manifests a binary distinction.

Some languages, such as French, have three parameters: Past, Present and Future.

Some languages have more than three parameters, distinguishing additionally remote past from recent past, for instance. The language with the most parameters thus far reported is an African language: Bamileke-Dschang with eleven. Crucially, parameters are encoded by specific lexical concepts, and thus form part of the knowledge bundle that constitutes a lexical concept. For instance, the parameter Past is encoded, in English, by the lexical concept associated with the *-ed* form in (3a). However, other lexical concepts also include the parameter Past such as the lexical concepts associated with the following forms: *sang, lost, went, etc.*

I argue, then, that a key feature of linguistic (as opposed) to conceptual content is that it encodes knowledge in parametric fashion. Parameterisation is a

highly reductive form of abstraction: it serves to abstract across the complexity exhibited by a particular category. In consequence the parameters encoded by linguistic content serves to ‘strip away’ most of the differences apparent in the original experience, thereby reducing it to a highly limited numbers of parameters.

4.2 The non-analogue nature of linguistic content

As conceptual content relates to records of perceptual states captured directly from a variety of experience types including sense perception, proprioception and subjective experience, it therefore consists, as noted above, of perceptual states recorded in analogue fashion: in a format that is similar to the perceptual experiences that gave rise to them. Indeed, there is a good deal of evidence, in the neuroscience literature, that sensory-motor representations, for example, are stored in the same areas of sensory-motor cortex that process sensory motor experience (Pulvermüller 1999).

In contrast, I argue that linguistic content is so highly schematic in nature that it is *non-analogue*: it takes a format that is not analogous to the perceptual experiences that it is a schematisation of. Hence, due to the reduction of rich perceptual information to highly impoverished parameters, this gives rise to a qualitatively very different type of information from the kind captured by conceptual content. To illustrate, re-consider the parameters Past and Non-past discussed with respect to example (3) above. These parameters are highly schematic abstractions drawn from the complex range of temporal relationships that hold between our experience of past, and our experience of now: our temporal location as experiencing centres of consciousness. Temporal experience, a form of subjective experience, is

extremely rich in perceptual terms (Evans 2004a). Yet the parameters Past and Non-past are not rich at all.

In sum, parameters encode highly schematic linguistic content abstracted from far richer perceptual experience, as recorded in the conceptual system, and provide a means for encoding recurrent ‘digitised’ dimensions of humanly relevant experience in an efficient way. In contrast, conceptual content which is accessed via open-class lexical concepts, gives rise to perceptually rich aspects of experience, about which I will have more to say later.

4.3 Topological reference

A further consequence of the highly reductive nature of the parameters encoded as linguistic content, and one first pointed to by Talmy (e.g. 2000), is that they provide *topological reference* rather than *Euclidean reference*. That is, linguistic content encodes schematic aspects of sensory-motor, proprioceptive and subjective experience, while conceptual content, to which open-class lexical concepts facilitates access, relates to precise, metric distinctions.

To illustrate consider the closed-class lexical concepts associated with the demonstrative vehicles *this* and *that*. The lexical concepts associated with these vehicles encode a distinction between an entity construed as proximal to the speaker, glossed as [THIS], versus an entity construed as distal, glossed as [THAT]. The distinction between the lexical concepts [THIS] versus [THAT] is illustrated by (4):

- (4) “Sit on this chair not that one!”

In this utterance, the chair that the addressee is being asked to sit on is the one closer to the speaker: ‘this chair’ as opposed to ‘that one’. Nevertheless, the distinction between [THIS] versus [THAT] does not rely upon precise metric details such as the exact distance from the speaker, in terms of metres, centimetres and millimetres, for instance. After all, it is immaterial how far the chairs are from the speaker (within reason), as long as one is closer to the speaker than the other. In other words, closed-class lexical concepts are *magnitude neutral*, where magnitude has to do with metric properties relating to distance. This is what it means to say that closed-class lexical concepts provide topological reference.

In contrast, the open-class lexical concepts, in addition to encoding linguistic content also, additionally, facilitate access to conceptual content, and hence can be employed to express metric details of distance giving rise to Euclidean reference. This is illustrated by (5):

(5) “Sit on the chair 2.54 metres away from me!”

The expression ‘2.54 metres’ involves open-class lexical concepts rather than closed-class lexical concepts, and serves to evoke with greater precision the chair in question.⁴

4.4 A restricted set of domains and categories

A consequence of parameterisation is that the range of *domains*, and the member *categories* that populate them, are highly restricted in terms of their encoding as

⁴ The parameters encoded as linguistic content exhibit a range of other *Euclidean neutralities*: notably with respect to the domains of SPACE and TIME. See Evans (To appear) for details.

parameters in linguistic content (cf. Talmy 2000). In using the term domain I have in mind large-scale and coherent bodies of knowledge such as the following: TIME, SPACE, COLOUR, MOTION, FORCE, TEMPERATURE, MENTAL STATES, and so on. By category I have in mind the member notions that populate a particular domain. For instance, in terms of the domain of TIME, categories consist of notions such as Punctuality, Durativity, Sequentiality, Simultaneity, Synchronicity, Boundedness, Time reference (e.g. Past, versus Non-past etc.), Time-reckoning (e.g. 10.05pm, etc.), and so forth. While *all* the domains of the sort just mentioned, and the categories which populate them, are evident at the conceptual level, only a restricted subset are encoded at the linguistic level, in terms of linguistic content.

For instance, some domains to which open-class lexical concepts facilitate access, such as COLOUR, do not appear at all in terms of linguistic content in English or any other language. That is, there are no parameters, in the sense defined above, that relate to this domain. This follows as many (perhaps most) domains do not relate to experience that can be straightforwardly parameterised in a humanly relevant way. There are at least two likely explanations for this. Firstly, the nature of the domain in question may not lend itself to being ‘reduced’ to highly schematised digitised parameters. After all, the reduction to content that does not directly give rise to simulations results in a reduction that, for some domains such as COLOUR, may eliminate the essential character of the information thereby making it uninterpretable. A second reason is that some domains do not relate in a ubiquitous way to the humanly relevant scenes that language serves to encode. For instance, categories that relate to the domain of MEDIAEVAL MUSICOLOGY, or even parameters that relate to less esoteric domains such as LOVE or JOURNEYS are not as ubiquitous in human

experience as parameters relating to domains such as SPACE, TIME, MOTION and MENTAL STATES which do appear to be encoded in linguistic content.

In addition to the restricted set of domains encoded, linguistic content also features only a small number of categories within each domain. To illustrate, consider a few of the categories associated with the domain TIME:

Domain: TIME

Category: Time-reference	Parameter
(6) a. He kicked the ball	Past
b. He kicks the ball	Non-past

Category: Boundedness	Parameter
(7) a. Holly has left the party	Bounded
b. Holly is leaving the party	Unbounded

Category: Plexity	Parameter
(8) a. Fred coughed	Uniplex
b. Fred coughed for 10 minutes	Multiplex

The category that I refer to as Time-reference is more traditionally referred to as tense. Of the other two categories illustrated, these are normally treated as relating to what is commonly referred to as aspect. The examples in (7) are usually referred to as perfective and imperfective aspect. The more usual terms for Uniplex and Multiplex in (8), as they relate to TIME, are 'semelfactive' and 'iterative' respectively.

4.5 Nominal versus relational lexical concepts

Another aspect of linguistic content is that it encodes a bifurcation between *nominals* and *relations* (Langacker 1987). The distinction in type of lexical concepts is as follows. Nominal lexical concepts are *conceptually autonomous*: they relate to entities which are independently identifiable, such as ‘chair’, or ‘shoe’. In contrast, relations are *conceptually dependent*: they constitute a relation holding between other entities, and are thus ‘dependent’ on those other entities in order to fully determine the nature of the relationship. For instance, in an utterance such as the following:

(9) Max hid the mobile telephone under the bed.

The lexical concept associated with the form *hid*, which I shall gloss as [HID], relates the conceptually autonomous lexical concepts associated with the vehicles *Max*, *mobile telephone* and *bed*, establishing a relationship involving ‘hiding’ between the conceptually autonomous participants in the linguistically-mediated *conception*⁵: namely [MAX] and [BED]. Analogously, the lexical concept associated with the vehicle *under* establishes a spatial relation between lexical concepts associated with *mobile telephone* and *bed*.

The conceptually dependent structure of relational lexical concepts is modelled, in LCCM Theory, in terms of a schematic *participant role* (Goldberg 1995). The lexical concept [HID] as exemplified in (9) encodes three schematic participant roles.⁶ The rich content relating to the participant roles is not specified in

⁵ ‘Conception’ is a technical term used in LCCM Theory to refer to utterance meaning.

⁶ Notice that the form *hid* is polysemous. For instance, *hid* is also associated with the ‘reflexive’ lexical concept in which an entity hides oneself, as in: *John hid in the wardrobe*. This lexical concept, which I gloss as [REFLEXIVE HID] encodes two schematic participant roles.

linguistic content. This arises from access to conceptual structure. That is, conceptual structure encodes rich content relating to hiding: that it involves someone that does the hiding for particular reasons, and that an entity of a particular sort, often an object, is hidden. Non-linguistic knowledge also includes what facilitates something being hidden, such as perceptual inaccessibility of the object being hidden and/or its being placed in a novel location. Conceptual content also includes information relating to the motor processes involved in hiding, which involve moving the object from one location to another. The participant roles encoded as part of the linguistic content for [HID] do not encode such details. Rather, what is encoded is a highly abstract representation, derived from the rich perceptual details of a hiding scenario. As such we have three roles that serve to distinguish between the three entities involved at the most general level of detail. These participant roles are: Hider, Object and Location.

Just as the bifurcation in lexical concepts discussed above — that holding between lexical concepts which solely encode linguistic content and those which additionally facilitate access to conceptual content — corresponds to a distinction in the formal encoding of lexical concepts: the distinction between open and closed-class vehicles, so too the distinction between nominal and relational lexical concepts has a formal reflex in terms of linguistic vehicles. In a language such as English, for instance, this distinction relates to lexical concepts associated with what are commonly referred to as nouns and noun phrases (nominals) on the one hand, and lexical concepts associated with other lexical forms, including verbs, prepositions, adjectives, adverbs and non-finite verb forms such as infinitives and participles (relations) on the other (see Langacker 1987 for details).

4.6 Referentiality

Another key aspect of linguistic content is that it is inherently referential in nature. Referentiality takes a number of different forms, as detailed below. However, the defining feature is that lexical concepts serve to encode the following: an intention that a particular entity is being indexed or, more informally, ‘pointed to’. In using the term ‘entity’ I have in mind physical entities that inhabit the world such as people, as well as physical artefacts, such as ‘Sam’ and ‘ball’ in (10a), abstract notions such as ideas, for example ‘peace’ in (10b), as well as relations that hold between physical entities and abstract ideas, such as ‘kicked’ in (10a) and ‘thought about’ in (10b), as well as highly schematic relations, as encoded by ‘to’ in (10c).

- (10) a. Sam kicked the ball
 b. Sam thought about peace
 c. Sam walked to the park

I identify at least three distinct types of *reference* encoded by lexical concepts.

The first type relates to what I will refer to as *denotational reference*. Many lexical concepts serve to index a physical entity of some sort, whether real or imagined. In this sense, part of what the lexical concepts associated with the vehicles *John* and *unicorn* serve to do is to signal an intention, on the part of the speaker, to refer to a given entity.

The second type I refer to as *cognitive reference*. This relates to relatively abstract notions or ideas that have no physical substance, whether real or imagined, and relate to lexical concepts associated with forms such as *love*, *war*, *phonology*, and

so forth. Hence, lexical concepts that serve to encode cognitive reference serve to signal an intention, on the part of the speaker, to refer to a non-physical idea.

The third type I refer to as *contextual reference*. This involves reference to an entity that is present in the linguistic or extra-linguistic discourse context. Hence, reference of this sort involves the encoding, by a lexical concept, of an intention to refer to an entity that the addressee can recover from context.

One type of contextual reference is *textual reference*. One form of textual reference involves reference to an entity already mentioned. This is traditionally termed *anaphora*. Textual reference that relates to an entity yet to be mentioned is termed *cataphora*. Examples of textual reference are provided in the examples below.

- (11) a. John is smart. **He** had a reading age of 14 by the time **he** was just 8.
 b. I want to say just **this**: I love you.
 c. The new target to reduce carbon emissions by 20% by 2020 will be a tough **thing** to achieve.

In the examples in (11a), the lexical concepts associated with the forms *he*, *this* and *thing* are specialised for referring to other entities (underlined) in the text.

There are many kinds of lexical concepts which encode an intention to signal contextual reference as it relates to extra-linguistic context. Many of these are often treated under the heading of *deixis*. Previous research has identified a range of diverse sorts of deictic lexical concepts including phenomena referred to as *spatial deixis*, *temporal deixis* and *social deixis* (for details see Fillmore 1997; Levinson 1983).

4.7 Pragmatic point

The final dimension of linguistic content that I address here relates to what I refer to as pragmatic point. This is a term I borrow from Fillmore et al. (1988). I use this term to refer to schematic aspects of the extra-linguistic dimensions of the encoding of linguistic content by a given lexical concept. As I use it, this term relates, broadly, to two aspects: i) the contexts of use in which a given lexical concept is conventionally employed, including settings and participants, and ii) some aspects of what has traditionally been referred to as the *illocutionary point* (Searle 1969) of a given lexical concept: which is to say the communicative purpose for which a lexical concept is employed.⁷

To illustrate the notion of pragmatic point consider the vehicle *declared* in the examples below. This is associated with at least three English lexical concepts, each of which exhibits a different pragmatic point.

- (12) a. She declared her love for him
- b. Chamberlain declared war on Germany on September 3rd 1939
- c. Despite being over the limit on the amount of dollars in cash eligible to be taken into the country, she declared nothing as she crossed the US border.

⁷ It is worth re-emphasising here that linguistic content is schematic in nature. Hence, while making a *speech act* (Searle 1969), such as declaring a state of war, for example, involves being able to call upon highly detailed bodies of conceptual knowledge relating to the sorts of scenarios and participants involved, linguistic content involves only the most generic aspects, including schematic information concerning the types of context in which a particular lexical concept can be deployed, the nature of the participants involved and the conditions which must hold.

The use of *declared* in (12a) serves to encode an intention to provide information of a particular sort, with an above-average level of assertiveness. Hence, the lexical concept which sanctions this use of *declared* can be glossed as [FORTHRIGHT INFORMATIONAL ASSERTION]. In contrast, the lexical concept associated with the use of *declared* in (12b) relates to an assertion which either changes, or otherwise revises, an institutional state. Crucially, not only is the illocutionary point distinct from the lexical concept responsible for the use of *declared* in (12a), but the context of use is distinct too. This follows as the context of use for the [ANNOUNCEMENT OF NEW LEGAL STATUS] in (12b) can only be successfully deployed by suitably qualified participants. For instance, Neville Chamberlain was able to successfully deploy this lexical concept because on September 3rd 1939 when he declared war, he was the legally-appointed Prime Minister of The United Kingdom, and under the terms of The Royal Prerogative — powers invested in the monarch and deployed by the Prime Minister on behalf of the monarch — he was legally entitled to take the country to war.

Finally, the lexical concept which sanctions the use of *declared* in (12c) relates to the [ANNOUNCEMENT OF DUTIABLE GOODS AT CUSTOMS] lexical concept. This is distinct both in terms of illocutionary point and context(s) of use from the previously mentioned lexical concepts. This lexical concept is specialised for use in contexts involving customs provision at international border crossings. Its communicative function has to do with signalling as to goods being transported, or caused to be transported by the person issuing the ‘declaration’ in this specific context, with respect to restrictions on the nature and/or amount of goods that may be transported into the country which establishes the customs provision, and/or tax payable on particular goods.

Based on the foregoing discussion, table 2 presents a summary of the key components of pragmatic point that are encoded as part of the linguistic content of each of the three lexical concepts. Much of the content associated with the three lexical concepts for declared comes from the conceptual content to which they afford access. However, pragmatic point, which concerns linguistic content, is highly schematic in nature. In these terms then, the distinction between the three lexical concepts relates to whether they stipulate that the setting is restricted or not, whether the participants are restricted or not, and the nature of the communicative function: the illocutionary point. Hence, by way of illustration, the lexical concept [ANNOUNCEMENT OF NEW LEGAL STATUS] encodes the following: there is no restriction on where the utterance can take place for it to realise its illocutionary point; the participants involved are, however, restricted, and the communicative purpose is to change some institutional state. This information is clearly highly schematic. However, it adequately captures, I argue, the highly stable aspects of the content associated with this lexical concept, which is to say, its linguistic content.

Table 2. Pragmatic point for three lexical concepts of *declared*

Lexical concept	Setting	Participant(s)	Illocutionary point
[FORTHRIGHT INFORMATIONAL ASSERTION]	Unrestricted	Unrestricted	Make statement
[ANNOUNCEMENT OF NEW LEGAL STATUS]	Unrestricted	Restricted	Change official state
[ANNOUNCEMENT OF DUTIABLE GOODS AT CUSTOMS]	Restricted	Restricted	Make official statement

5. The nature of conceptual structure

In this section I am concerned, in broad terms, with conceptual structure: the nature and organisation of concepts. In LCCM Theory conceptual structure is modelled in terms of the cognitive model. A cognitive model is, in essence, similar to Barsalou's (1999) notion of a *simulator*.⁸

A simulator (Barsalou 1999) constitutes records of perceptual states, stored in a coherent format referred to as a *frame*, which can be re-activated, often in novel ways. The re-activations are referred to as *simulations*, (e.g. Barsalou 1999, 2003; Gallese and Lakoff 2005; Kaschak and Glenberg 2000; Prinz 2002; Glenberg and Kaschak 2002; Zwaan 1999, 2004). The mechanism known as *simulation* represents a general purpose computation performed by the conceptual system in order to recover the bodily states stored within frames and to perform operations deploying such perceptual states. As such, a frame can give rise to a potentially limitless set of re-activations or simulations. Hence, a simulator, and thus a cognitive model, encompasses a frame and a potentially unlimited set of simulations.

⁸ The use of a novel term, 'cognitive model', is done for two reasons. Firstly, at this stage in our understanding, it is not clear to what extent units of semantic structure: lexical concepts, facilitate access to the conceptual system. For instance, the common experience of 'not being able to put thoughts into words', particularly as applied to subjective experiences, suggests that the linguistic system may be less well connected to certain types of conceptual representations than others. Indeed, this is a point made by Jackendoff (e.g., 1992). It is conceivable that some aspects of conceptual structure may only be partially accessible or even inaccessible to the linguistic system. I introduce the theoretical construct of the cognitive model, then, to distinguish between those simulators which are accessible via linguistic representations, and those which are not. Simply put, while the conceptual system is populated by simulators (Barsalou 1999), cognitive models are simulators which are specialised for being accessed by lexical concepts. Hence, the rationale for introducing the term 'cognitive model' is to identify those simulators which with which the linguistic system interacts. The second reason is as follows. In his development of the notion of a simulator, Barsalou is primarily focused on the perceptual basis — in the wider sense as described earlier — of conceptual structure. While he acknowledges that other forms of information are likely to feed into conceptual representations, he is primarily exercised by accounting for the perceptual grounding of cognition. In my account, I explicitly acknowledge that propositional (i.e., non-perceptual) information may also become incorporated in cognitive models, which supplements the perceptual information already present. Such propositional information is likely to accrue via linguistically mediated routes, including narrative, exchange of news, and gossip. For these reasons, it is useful to distinguish the theoretical construct under development here, by applying the novel term cognitive model.

Following Barsalou, and indeed others who take an *embodied* or *grounded cognition* perspective (see Barsalou 2008 for a review), I assume that the perceptual states that make up cognitive models derive from a number of sources, as briefly introduced earlier. These include: i) the processing of external stimuli via sensory (or modal) systems (vision, audition, olfaction, haptics, and gustation); ii) action, which provides motor information relating to bodily states via proprioception: information about movements involving joints and muscles, as well as the vestibular system, which provides information as to position in space and motion trajectories.⁹ In addition, subjective (or introspective) experiences are just as important for giving rise to records of perceptual states that make up cognitive models. For instance, Damasio (e.g. 1994) in ground-breaking work on emotion has emphasised a number of categories of feelings that arise from internal body states. These include body states (emotions) that we label as Happiness, Sadness, Anger, Fear, and Disgust. These give rise to phenomenologically real, in the sense of directly experienced *feelings*. Damasio identifies a further category of feeling, what he terms *background feelings*, which derive from internal body states. Background feelings arise from, among other things, interoceptive experience, which is to say the visceral sense — our felt sense of the internal organs and other internal bodily states. Other subjective experiences, which are directly felt, include various aspects of temporal experience which arise from bodily states (circadian rhythms such as the wake-sleep cycle), as well as perceptual processing, which is subserved by a wide range of neurologically instantiated temporal mechanisms (see Evans 2004a, 2004b and references therein), and consciousness (Chafe 1994; Grady 1997). Hence, the perceptual states that make up the frames and give rise to the simulations that comprise given cognitive models

⁹ See Evans (To appear) for a review of the operation of the sensory mechanisms responsible for sense-perception.

are grounded in both sensory-motor experience and subjective experience: experience of internal bodily and cognitive states, including emotion, mood and affect.

5.1 Perceptual symbols

Following Barsalou (1999) I assume that individual records of perceptual states are stored as *perceptual symbols*. It is well known from research on attention that during perceptual experience, the cognitive system can focus attention on individual components of the stimulus array. For instance, attention can selectively focus on the colour of an object, filtering out, for instance, its shape, or texture, and even the surrounding objects (Garner 1974, 1978). Through selective attention, individual perceptual components derived from perceptual experience of the kinds discussed above are recorded, in bottom-up fashion, in sensory-motor areas of the brain (Barsalou 1999).¹⁰ The components are stored in schematic fashion. This means that it is not individual perceptual states that are stored, but rather commonalities are abstracted across specific instances of perceptual states providing individual memories deriving from sense perception (e.g., individual memories for *red*, *hot* and *purr*), proprioception (e.g., *lift*, *run*) and subjective experience (e.g., *compare*, *similar*, *hungry*). These schematic memories Barsalou refers to as perceptual symbols. They are symbols in the sense that, later, in top-down fashion, they can be reactivated, or simulated, and can be used to support the range of symbolic behaviours that subserve a fully functional conceptual system.

¹⁰ There is compelling neuropsychological and neuroimaging evidence which supports the view that human conceptual representations are grounded in the modalities, and hence are perceptual in nature. For instance, categorical knowledge is grounded in sensory-motor regions of the brain (for reviews see Damasio 1989; Pulvermüller 1999, 2001). Damage to a particular sensory-motor region serves to impair the processing of categories that use the region in question to perceive physical exemplars.

Perceptual symbols implement a conceptual system as follows. Barsalou argues that memories of similar and related components become organised into a system of perceptual symbols which exhibit coherence: a *perceptual symbol system* (Barsalou 1999). This perceptual symbol system is what I refer to as a frame. A frame, then, is an information structure consisting of large collections of perceptual symbols, encoding information which is stable over time as well as incorporating variability. Hence, a frame provides a unified, and hence coherent, representation of a particular entity. For instance, a frame involves numerous components that have a perceptual basis, that are related in various ways. In addition, the perceptual symbols that collectively comprise the frame can be combined in a range of ways, giving rise to an infinite variety of simulations. Hence, a system of perceptual symbols gives rise to both a frame: a relatively stable knowledge matrix, and dynamic simulations.

5.2 Frames

In this section I identify a number of frame types. I do so based on Barsalou's work on frames (e.g. Barsalou 1991, 1992), and Barsalou et al. (1993). In broad terms, frames can be identified which relate to *things* and to *situations*. Further, within each of these broad divisions there are frames which are *episodic*, relating to specific types of experience and/ or knowledge and frames which are *generic*, relating to schematisation over broadly similar aspects of experience and/ or knowledge. The distinct frames identified below are *individuals* (episodic) and *types* (generic), which relate to things, and *episodic situations* and *generic situations*, which relate, self-evidently, to situations. I begin by focusing on the frames for things: individuals and types, before proceeding with a discussion of the frames for situations.

THE WORLD MODEL

Barsalou et al. (1993) provide an ontology for a theory of knowledge representation, which is based on what they refer to as the *world model*. This comprises a person's beliefs about the current state of the world. These beliefs relate to individuals, their current states and where they are located. Barsalou suggests that people employ a hierarchically-arranged core of spatial frames. That is, people represent the world and its contents in a spatial fashion, corresponding to continents, countries, cities, neighbourhoods, individual buildings, rooms and locations within rooms. They further locate entities within these locations, and integrate the spatial frames with temporal knowledge, for instance, relating to cycles and time-frames of various sorts including the seasons, the calendar, and temporal intervals such as years, months, weeks and days, as well as content-based temporal structures such as knowledge relating to one's own and family members daily routine, development over the life span, stages in career progression, and so on. Temporal information serves to organise past, present and future information in the world model and, Barsalou argues, does so orthogonally to the spatial core. Moreover, in this world model, people represent people's interactions and movements, updating the model continuously. For instance, while at work, a person might represent their partner's movements, going to the shops, returning home, or their children's activities while at school, and so on. People also represent other ongoing activities taking place in the various regions represented in their world model. For instance, one might know about a meeting of a University Exam Board taking place in a committee room near one's office, it being Tuesday afternoon, Prime Minister's question time taking place at the House of Commons, knowing — based on having read today's newspaper — that the Queen is currently staying at Windsor Castle rather than Buckingham Palace, that Big

Ben in London is currently undergoing repairs and hence not presently chiming, and so on.

In the world model, two distinct kinds of frames can be distinguished which relate to things: individuals and types (Barsalou et al. 1993).¹¹ Individuals are frames that relate to animate and inanimate entities that are held to persist continuously in the environment. As such, individuals are central to the ontology of the world model. Individuals provide relatively stable information about a given entity: information that is both stable over time, as well as incorporating episodic information. Hence, the new information for a given individual is added to the frame thereby updating it on an ongoing basis. An individual is updated based on encounters with the entity it represents. For instance, the frame for ‘my car’ might include the petrol gauge reading the last time I interacted with it, and the fact that I have noticed there is an oil leak, and that the car needs cleaning. This information is merged into the frame to provide an updated representation.

Crucially, although the same individual may be encountered in the world on many occasions, often in the same day, in terms of the world model all the episodic information extracted during these encounters is integrated into the individual frame. This follows from the *one-entity one-frame principle* (Barsalou et al. 1993). This principle holds that only one frame can relate to any given entity. Hence, all the information extracted from experience, which is related to a particular individual, is merged into the frame for that entity. Hence, the frame for a particular colleague at work may include information relating to his location the last time I interacted with him, and so on.

¹¹ Barsalou et al. (1993) use the term ‘model’ to refer to what I am here calling ‘type’. I prefer the more intuitively accessible term ‘type’ and also seek to avoid any confusion with the construct of the cognitive model. Hence, I do not use the term ‘model’.

In addition to individuals, Barsalou et al. (1993) argue that there is another frame type which inheres in the world model. This type of frame, which I refer to as: type, is an abstraction across frames for individuals providing a frame for a type of individual. As such, types are not conceptualised as having corresponding entities in the world. For example, while the individual for 'my car' in the world model corresponds to my car in the world, the frame for 'car' is a type, and relates to a type of individual, abstracted from across a range of individuals. Hence, people understand their frames for types to inhere only in the world model, but not, crucially, in the world itself.

One of the features of individuals in the world is that they change location. In the world model, this feature is captured in terms of the phenomenon referred to as *transcendence* (Barsalou et al. 1993). Transcendence has to do with the number and range of locations at which individuals and types are represented. For instance, a colleague from work will be represented at work. However, a chance meeting at the local supermarket will ensure that the individual frame for the colleague becomes additionally stored at the supermarket location in the world model. When the colleague goes on vacation to Paris, and sends a postcard in to the office to report on the vacation, the individual is additionally stored as part of the Paris location in the world model.

Barsalou et al. (1993) argue that transcendent frames for individuals and types, while being located at multiple sites in the world model, become *functionally detached* from the world model. That is, they give rise to a level of information about the nature of individuals and types, and the interactions they can engage in which become abstracted from the spatial frames that form the core of the world model. In other words, transcendence gives rise to de-contextualised representations which form

transcendent taxonomies. For example, the type for ‘heart’ is a feature of all mammals. Hence, its presence as part of the frame for numerous individuals and types gives rise to transcendence.

This property serves two important functions. Firstly, transcendence provides an important means of organising beliefs about the nature of entities in the world. It does so as it serves to capture similarities between individuals and models. As such, it facilitates inferences. For instance, we can infer that lions have hearts on the basis of knowing that all mammals possess hearts. Secondly, transcendent taxonomies may constitute important building blocks in the construction of the world model. This follows as transcendent information can be inserted into frames for new individuals upon first encounter. For instance, on encountering an unfamiliar cat, information from the model for cats is retrieved and copied, in order to form the basis for the new individual in the world model. This process serves to minimise the amount of learning about new entities before they can be adequately represented.

SITUATIONS

Having briefly described the ontology for individuals and types, I now consider how situations are modelled. The basic insight is that in addition to individuals and types, humans additionally represent situations, there being two kinds of situation: episodic situations and generic situations. The distinction between episodic and generic situations is orthogonal to the distinction between individuals and types.

According to this approach, situations are part of larger *events* — events are composed of situations — while being made up of discrete *images*. As with situations, events and images, as I use the terms, are mental representations. The notions event, situation and image are somewhat akin to the notions of scripts, scenes

and states developed in Schank (1975, 1982), and Schank and Abelson (1977), with the difference being that events, situations and images are made up of perceptual symbols, and hence are perceptual and thus embodied in nature.

One of the key insights of this approach is that it takes a *situated cognition* perspective. That is, people's frames for individuals and types are *situated*, and *local* rather than being de-contextualised and universal. An individual or type is situated in the sense that it is represented in the situations in which it occurs. For instance, the individual frame for 'my sofa' is represented as being located in my living room. Hence, the frame for my sofa is related to the situation frame for 'my living room'. Similarly, individuals and types are local in the sense that they relate only to exemplars actually encountered, rather than being generalised to entities universally. For instance, the type for 'sofa' incorporates information relating only to sofas that have been encountered. In this way, this approach to knowledge representation assumes that the conceptual system is directly grounded in situated action and interaction.

Barsalou et al. (1993) propose that the mental representations they refer to as images are static spatial scenes (cf. Tyler and Evans 2003). These may consist of frames for individuals and/or types, viewed from a particular viewpoint, with a particular geometric, topological and functional relationship holding between them. Crucially, an image is composed of numerous perceptual symbols. For instance, a person may represent a picture hanging on the wall above the sofa in their living room.

A situation is comprised of a series of images. Hence, and as with an image, a situation may consist of a relatively stable set of individuals and types. The difference is that a situation, while occupying a relatively constant region of space is

dynamic, in the sense that entities may interact and move around, and there is change over time. For instance, a situation might involve a person approaching the sofa, sitting down, turning their head to look at the picture on the wall, turn their head away again, sitting for a while, before getting up and moving away from the sofa.

An event comprises a series of two or more situations which are related in coherent fashion. The key difference between an event and a situation is that an event involves a significant outcome, often involving a change in regions of space and/or the individuals and/or types involved in the event. For instance, an event might involve a person going to a department store and purchasing a picture, bringing it home in their car, fetching a hammer and nail from the garage, selecting a spot on the wall above the sofa to hang the picture, knocking a nail in the wall at the desired location, and hanging the picture above the sofa. A table summarising the differences between image, situation and event *qua* mental representations is provided in table 3.

Table 3. Features of images, situations and events

Features of Images	Features of situations	Features of events
i) a set of perceptual symbols ii) represents individuals and/or types iii) a static spatial configuration iv) viewed from a particular perspective	i) a series of images ii) depicts a relatively constant set of individuals and/or types iii) depicts some significant change over time iv) occurs in a relatively constant region of space	i) a series of two or more situations ii) the situations are related in a coherent manner iii) the situations lead to a significant outcome

As observed above, there are distinct sorts of frames relating to both episodic and generic situations, which parallels the distinction between individuals and types. An

episodic situation arises from perceiving a situation in the world, the situation *qua* frame constituting a mental representation of the perceived situation. Moreover, humans represent situations at the locations in their world model where the situation occurs. For instance, in the example of the situation involving the hanging of a picture above the sofa, the frame for the episodic situation is linked to the frame for the conceptualiser's living room. On this account, and just as we saw with frames for individuals above, episodic situations are not wholly episodic. They also include a potentially large amount of generic information. This is due to the phenomenon of transcendence, which facilitates cognitive economy: generic knowledge can be shared between related frames. As with frames for things: individuals and types, discussed above, frames for situations are associated with temporal knowledge structures such as those relating to daily routines, life periods, hours of the day, and so on.

In contrast, frames for generic situations do not include episodic information. Rather they develop by virtue of abstracting away points of difference, in order to distil the commonalities that persist in different frames for episodic situations. Like frames for types, discussed above, generic situations do not have direct counterparts in the world.

Barsalou et al. (1993) propose that frames for a generic situation are formed when two or more episodic situations share a number of commonalities. These are presented in table 4. These commonalties serve to indicate that two episodic situations are related. The episodic situations in question are then abstracted in order to form a generic situation for this type of situation.

Table 4. Identification of commonalities in the formation of an abstract situation (after Barsalou et al. 1993)

Two situations are related when the following occur:
<ul style="list-style-type: none"> i) They share a common number of images. ii) They share common individuals and/or types. iii) The configuration of individuals/types in each similar image across situations is qualitatively the same. iv) The transformations of individuals/types between similar images across situations is qualitatively the same. v) The two situations culminate in a common end state.

5.3 The structure of frames

Frames have three basic constituents: *attribute-value sets*, *structural invariants* and *constraints*. In this section, which draws on Barsalou (1992) I examine each of these in turn.

ATTRIBUTE-VALUE SETS

Frames consist of sets of attributes and values. An attribute concerns some aspect of a given frame, while a value is the specification of that aspect. For example, in terms of the vastly simplified frame for CAR depicted in figure 4, ENGINE represents one aspect of the CAR, as do DRIVER, FUEL, TRANSMISSION and WHEELS. An attribute is therefore a concept that represents one aspect of a larger whole. Attributes are represented in figure 4 as ovals. Values are subordinate concepts, which represent subtypes of an attribute. For instance, SUE and MIKE are types of DRIVER; PETROL and DIESEL are types of FUEL; MANUAL and AUTOMATIC are types of TRANSMISSION, and so on.

Values are represented as dotted rectangles in figure 4. Crucially, while values are more specific than attributes, a value can also be an attribute, because it can also have subtypes. For instance, PETROL is an attribute to the more specific concepts UNLEADED PETROL and LEADED PETROL, which are values of PETROL. Attributes and values are therefore super-ordinate and subordinate concepts within an *attribute taxonomy*: subordinate concepts, or values, which are more specific, inherit properties from the super-ordinate concepts, or attributes, which are more general.

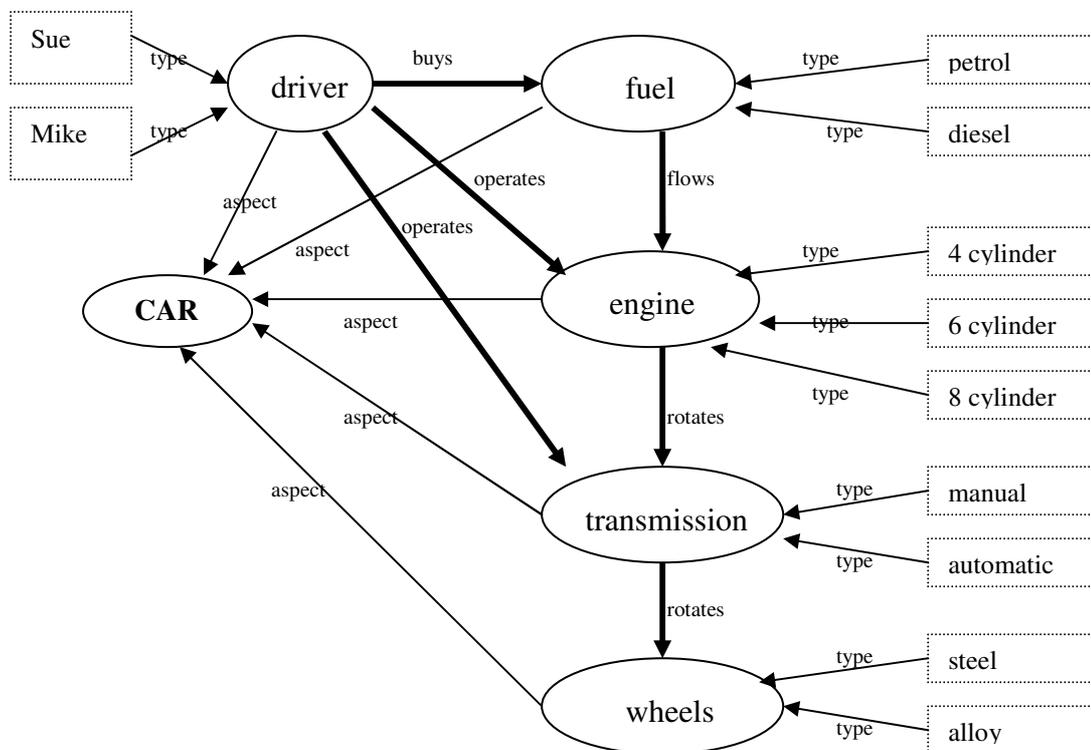


Figure 4. Frame for CAR (adapted from Barsalou 1992: 30).

In addition, attributes within a frame can be associated with their own *attribute frame*, providing an embedded form of framing. For instance, the attribute DRIVER in the CAR frame may have a number of attributes associated with it, including AGE, SEX, STATUS OF DRIVING LICENCE (i.e. whether it is 'clean' or not), NUMBER OF YEARS EXPERIENCE,

and so on. As frames are dynamic entities, undergoing continuous updating, attributes can be added to frames based on new encounters, or in order to achieve a particular goal. For instance, in the light of the recent introduction of a new banding scheme for road tax — an annual tax paid on all vehicles in the UK to use the public highway — based on petrol consumption, UK car owners are likely to have added a new attribute to their frame type for CAR relating to CAR TAX LEVEL. It is also worth emphasising that attribute-value sets, as with other aspects of knowledge representation, are likely to be idiosyncratic, and hence to vary from person to person.

A final property of attribute-value sets that I mention relates to what Barsalou and Billman (1989) have referred to as *attribute systematicity*. This concerns the idea that certain attributes are core, in the sense that they frequently recur across contexts.¹² This can facilitate frame formation. For instance, if a particular value for an attribute is not known when setting up a new frame of the type individual, a value for a core attribute can be ascribed based on the core attribute set retrieved from memory. For instance, imagine your friend is proudly showing off his new bright red sports car to you, a core attribute of the type frame: SPORTS CAR, is FUEL with the value PETROL. Hence, even though there may be no direct evidence that the car takes fuel, for instance, because you haven't noticed a petrol cap, or seen evidence of a fuel tank, this is something that will be added to the frame for this individual, and the value PETROL will be added as a consequence.

STRUCTURAL INVARIANTS

According to Barsalou, “[A]ttributes in a frame are not independent slots but are often related correlationally and conceptually” (Barsalou 1992: 35). In other words,

¹² As is well known, correlations in experience give rise to associative strength in memory: co-occurrence gives rise to a core set of attributes, which thus exhibit systematicity. See references in Barsalou et al. (1993) for instance.

attributes within a frame are related to one another in consistent ways across *exemplars*: instances of a given frame in the world. For example, in most exemplars of the frame CAR it is the driver who controls the speed of the ENGINE. This relation holds across most instances of cars, irrespective of the values involved, and is therefore represented in the frame as a *structural invariant*: a more or less invariant relation between attributes DRIVER and ENGINE. In figure 4 structural invariants are indicated by bold arrows. Hence, a structural invariant constitutes what Barsalou (1992) terms ‘a normative truth’ holding between attributes within a frame.

CONSTRAINTS AND FACTORS

Like structural invariants, constraints and factors are relations that hold between attributes, or more specifically, between attribute values. However, rather than capturing normative relations, constraints and factors give rise to variability in the values associated with attributes. This follows as values in a given frame are interdependent on the values associated with other attributes. There are two kinds of constraints, which I briefly review below, and two factors. The constraints are *global constraints*, and *local constraints*. The two factors are *contextual factors* and *goal factors*.

Global constraints serve to constrain attribute values globally. This means that a modification in one value entails a proportional modification in a related value. For instance, consider the example of a TRANSPORTATION frame involving a passenger in a taxi, for instance, being transported from one location to another. In this frame there is a negative attribute constraint which holds between the attributes SPEED and DURATION. That is, as the value for the attribute SPEED increases (and transportation becomes faster), so the value for the attribute DURATION decreases.

Local constraints constrain sets of values locally, rather than globally. That is, the presence of a given value entails the presence of a related value, while the absence of one entails the absence of another. For instance, consider a frame for VACATION. If the attribute ACTIVITY has the value SKIING, then this requires that the attribute HOLIDAY DESTINATION has the value SKI RESORT. Similarly, if the attribute ACTIVITY has the value SURFING, then the destination attribute must have the value OCEAN BEACH.

Contextual factors relate to aspects of context which serve to influence attribute values. For instance, the activity of SKIING requires a SKI RESORT, while increasing SPEED of travel reduces the DURATION of the journey. As aspects of situations are related rather than being independent, context constitutes a factor which can influence both global and local constraints.

Now I consider goal factors. In addition to context, an agent's goal(s) also provides a factor that influences the interaction between values associated with related attributes. For instance, in a PHYSICAL WORK OUT frame, the agent's goal, to get fit, serves to ensure that the attribute EXERTION forms part of the frame.

5.4 Chaining within the conceptual system

In this section I briefly consider the phenomenon of *chaining* (Barsalou et al. 1993; see also Lakoff 1987). The conceptual system is not a haphazard collection of cognitive models. Rather, cognitive models exhibit a range of often complex interconnections. As such, cognitive models are linked in a web of interconnections, of diverse sorts: hence, chaining. The consequence of this, in terms of linguistic

interaction, is that access sites established by lexical concepts provide a deep semantic potential for purposes of linguistically-mediated communication.

Chaining is a consequence of a number of different types of interconnections and relationships holding between frames. One such interconnection arises due to the phenomenon of attribute frames, discussed above. That is, frames are embedded within larger frames. Take the frame CAR, discussed above. A salient attribute associated with this type is ENGINE. The knowledge of engines possessed by one group of human conceptualisers, namely car mechanics, is highly complex, and this attribute includes many subordinate attributes each with corresponding values, which are themselves subordinate attributes with further values, and so on. In this way, a frame subsumes multiples frames which are embedded, capturing aspects of the larger units of which they are subparts.

Another way in which chaining occurs arises from the phenomenon of transcendence. This relates to the situated nature of cognitive models for things: individuals and types. Recall that cognitive models of this kind are 'located' in situations. In other words, cognitive models for things are located in the world model at the points at which they are encountered. Hence, cognitive models for episodic and generic situations include representations for individuals and types. The greater the number of situations to which individuals and types are linked the greater their transcendence is held to be. Hence, transcendence is a function of how interconnected cognitive models for things are with the range of representations for situations, and hence events, with which they are connected.

Another motivation for chaining arises due to the componential nature of the conceptual system itself. Recall that cognitive models are comprised of sets of perceptual symbols. As perceptual symbols are records of discrete perceptual states

(e.g. *purr*, *red*, *hot*, etc.), similar perceptual symbols (e.g. *red*) form part of many different cognitive models within the conceptual system. As such, unique records of similar perceptual states persist throughout the conceptual system. The consequence of this is that the conceptual system is thorough-goingly redundant in terms of the nature of the representations which make up the range of cognitive models which populate it. This provides, naturally, commonalities across cognitive models, and is a consequence of a fundamental design feature of the conceptual system.

Another way in which chaining arises is due to the relationships that exist between cognitive models, due to, broadly, the distinction between episodic versus generic cognitive models. For instance, in terms of cognitive models for things, we have the distinction individuals and types. While individuals may be related to each other based on the dimensions of chaining mentioned in the preceding paragraphs, a type is related to all the individuals from which it is formed. Similarly, a generic situation is related to all the episodic situations that it resembles, and from which it has abstracted across to provide a generic situation.

6. Interaction between the linguistic and conceptual systems

A key feature of knowledge representation in humans is that the linguistic system interacts with the conceptual system in order to facilitate access to conceptual knowledge. Indeed, as the philosopher of science Jesse Prinz (2002: 14) has observed:

Concepts must be capable of being shared by different individuals and by one individual at different times. This requirement...must be satisfied if concepts are to play some of their most important explanatory roles...it is almost universally assumed that concepts play a pivotal role in linguistic communication.

Indeed, a fundamental design feature of human cognition is that linguistic representations provide an indexing and control function, greatly increasing the range of uses and flexibility of the human conceptual system. However, this does mean that linguistic representations are equivalent to the concepts which populate the conceptual system.

I assume that the human conceptual system is, *en grandes lignes*, essentially the same as the primate conceptual system. Recent findings suggest that such an assumption is not unreasonable (e.g. Barsalou 2005; Hurford 2007). Given the relatively recent emergence of language, and the far greater antiquity of the conceptual system¹³ I assume that linguistic representations evolved to complement and enhance the existing form of representations that inhere in the conceptual system, rather than duplicating them.

From the perspective of LCCM Theory, the interaction between the linguistic and conceptual systems is facilitated by what I earlier referred to as open-class lexical concepts. I discuss the nature of the interaction by examining some of the relevant issues below.

¹³ For discussion, a sample of relevant book length treatments from various perspectives include Corballis (2002), Deacon (1997), Donald (1991), Dunbar (1996), Mithen (1996), Hurford (2007). See also the excellent collection of papers in Christiansen and Kirby (2003).

6.1 Access sites

The primary way in which the representations inhering in the linguistic and conceptual systems interact is by virtue of *access sites*. An access site is a theoretical construct in LCCM Theory which represents a composite of the range of *association areas* that hold between an open-class lexical concept and the conceptual system. An association area is a location in the conceptual system with which a specific lexical concept is associated. In other words, an association area provides a point of convergence between the two systems facilitating interaction between content from both. As a given lexical concept has typically many association areas, an access site constitutes the set of association areas for a given lexical concept. For example, and as we shall see below, the lexical concept [RED] is associated with many representations for individuals and types, each with its own distinctive hue, throughout the conceptual system. All the association areas collectively comprise the access site for this lexical concept. Yet the complexity of the way in which [RED] facilitates access to conceptual structure gives rise, as we shall see, to a large *semantic potential*.

The purpose of an access site is to facilitate integration of linguistic and conceptual content in order to provide an *integrated simulation*.¹⁴ Hence, the evolutionary motivation, on this account, for the linguistic and conceptual systems to interact is in order to make use of conceptual structure inhering in the conceptual system in service of linguistically-mediated communication.

I hypothesise that the association areas that comprise an access site arise by virtue of usage patterns: vehicles sanctioned by specific lexical concepts being used in

¹⁴ An integrated simulation is equivalent to what I what I have referred to as a conception: the meaning derived from compositional processes involved in understanding a well-formed utterance.

the context of perceived things and situations. Based on such patterns of use, statistical frequencies are extracted which serve to associate lexical concepts with the regions of the conceptual system where such things and situations are represented, giving rise to association areas. Access sites are thus probabilistic, in the sense that the greater the frequency with which a language user experiences a sanctioning lexical concept and a thing/situation as co-occurring, the greater the strength of the association area.¹⁵

6.2 Semantic potential

One consequence of the chaining exhibited by the conceptual system is that lexical concepts, by encoding access sites, facilitate access to a large semantic potential. To illustrate, let's briefly re-consider the lexical concept [RED] associated with the form *red*. The lexical concept [RED] facilitates access to a bewildering number of distinct perceptual symbols which contribute to a vast number of cognitive models in the conceptual system of any language user of English. To get a sense of the semantic potential involved, consider all the individuals and types that a single person will represent in their world model that features the perceptual state I gloss as *red*.

Limiting ourselves to types we might list Royal Mail post boxes, red squirrels, foxes, roses, blood, lipstick, Santa Claus' clothes, a robin's throat, strawberries, the red stop sign on the public highway, tomatoes, red traffic light, red cross, the flag of St. George, celebrity carpets, Babybel cheese wax, chilli peppers, fire engines, the Chinese flag, red wine, superman's cape, fire, henna, and so on. Notice that the represented hue associated with these types may vary from person to person, based on

¹⁵ See Barsalou et al. (To appear) for discussion of a related proposal. See also Boroditsky and Prinz (To appear).

cultural experience, and so on. Nevertheless, we can imagine contexts in which we would apply the phonological vehicle *red* in order to evoke the colour associated with these types.

In addition, there are further situations, both episodic and generic, that involve the individuals and types which include a perceptual symbol that I gloss as *red*. However, each of these perceptual symbols is unique to the individual and/or type and hence the situation of which it forms part. After all, it is the generic situation in which a teacher scrawls red ink on a pupil's exercise book, evoking a different perceptual symbol than the one evoked when we simulate a red squirrel scurrying up a tree. Nevertheless, the lexical concept [RED] is associated with, and hence facilitates access to, both. Put another way, the semantic potential for the lexical concept [RED] comes from the diverse range of perceptual symbols that are found in these cognitive models, and many others. Moreover, it is precisely because [RED] facilitates access to such a diverse potential that the vehicle *red* exhibits such variation in the way it can be used, as exhibited by the very different simulations we achieve for 'red' in the examples discussed earlier in section 3: the red associated with a red squirrel versus the red ink of a school teacher's pen.

6.3 The uniqueness of the access site

While lexical concepts are typically associated with a number (often many) cognitive models, which thereby make up the access site, the exact nature of the access site with which a lexical concept is associated is held to be unique. Put another way, no two lexical concepts share the same access site. While the range of cognitive models to which lexical concepts may be similar, they may never be exactly the same. The

consequence of this is that each lexical concept has a unique *cognitive model profile*: the range of cognitive models which make up an access site,. From the perspective of the linguistic system, this means that there can be no true synonymy between lexical concepts.

To illustrate, consider the lexical concepts which I gloss as [SHORE] and [COAST] associated with the forms *shore* and *coast*, respectively. As observed by Fillmore (1982) while the semantic representation for these two lexical concepts is similar it is not identical. This follows, in present terms, as while each of these lexical concepts exhibits partial overlap in the primary cognitive models, there are also distinctions. For instance, both lexical concepts facilitate access to a cognitive model profile relating to the strip of land that borders land and sea. However, each lexical concept accesses a cognitive model relating to a generic situation from which this land region is viewed. In the case of [SHORE] this concerns a sea-based perspective, i.e., on board a ship. In contrast, [COAST] does so from the perspective of land-based location. For this reason, a shore-to-shore trip is across water while a coast-to-coast trip is over land.

7. Summary

This paper has been concerned with developing an account of semantic representation, as assumed by LCCM Theory. LCCM Theory assumes a principled separation between the evolutionarily earlier conceptual system and the more recent linguistic system. Each system is populated by different types of ‘semantic’ representation: the lexical concept and the cognitive model. Moreover, the nature of

the content associated with the two systems is of a fundamentally different type. Linguistic content, encoded by lexical concepts, is highly schematic in nature, providing a structuring function to simulations. In contrast, conceptual content, encoded by cognitive models, provides perceptually rich and highly detailed information. In addition to encoding linguistic content, a subset of lexical concepts — open-class lexical concepts, serve as access sites, thereby facilitating interaction between linguistic and conceptual content, thereby giving rise to integrated simulations. In essence, LCCM Theory assumes that the linguistic system provides an executive control function, allowing access to conceptual representations for purposes of linguistically-mediated communication.

References

- Barsalou, L. 1991. Deriving categories to achieve goals. In *The Psychology of Learning and Motivation: Advances in Research and Theory*, Volume 27, G.H. Bower (ed.), 1-64. San Diego, CA: Academic Press.
- Barsalou, L. 1992. Frames, concepts, and conceptual fields. In *Frames, Fields, and Contrasts: New Essays in Lexical and Semantic Organization*, A. Lehrer and E.F. Kittay (eds), 21-74. Hillsdale, NJ: Lawrence Erlbaum.
- Barsalou, L. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22: 577–660
- Barsalou, L. 2003. Situated simulation in the human conceptual system. *Language and Cognitive Processes* 18: 513-562.

- Barsalou, L. 2005. Continuity of the conceptual system across species. *Trends in Cognitive Sciences* 9: 309-311.
- Barsalou, L. 2008. Grounded cognition. *Annual Review of Psychology* 59: 617-645.
- Barsalou, L. and Billman, D. 1989. Systematicity and semantic ambiguity. In *Resolving Semantic Ambiguity*, D. Gorfein (ed.), 146-203. New York: Springer-Verlag.
- Barsalou, L., Yeh, W., Luka, B., Olseth, K., Mix, K. and Wu, L.-L. 1993. Concepts and meaning. In *Chicago Linguistics Society 9: Papers from the Parasessions on Conceptual Representations*, Volume 2, K. Beals, G. Cooke, D. Kathman, K.E. McCullough, S. Kita and D. Testen (eds), 23-61. Chicago Linguistics Society.
- Barsalou, L., Santos, A., Simmons, K. and Wilson, C.D. To appear. Language and simulation in conceptual processing. In *Symbols, Embodiment, and Meaning*, M. De Vega, A. Glenberg and A. Graesser (eds). Oxford: Oxford University Press.
- Boroditsky, L. and Prinz, J. To appear. What thoughts are made of. In *Embodied Grounding: Social, Cognitive, Affective, and Neuroscientific Approaches*, G. Semin and E. Smith (eds). Cambridge: Cambridge University Press.
- Chafe, W. 1994. *Discourse, Consciousness, and Time: The Flow and Displacement of Conscious Experience in Speaking and Writing*. Chicago: University of Chicago Press.
- Christiansen, M.H. and Kirby, S. 2003. *Language Evolution*. Oxford: Oxford University Press.
- Corballis, M. 2002. *From Hand to Mouth: The Origins of Language*. Princeton: Princeton University Press.

- Damasio, A. 1989. Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33: 25–62.
- Damasio, A. 1994. *Descartes' Error: Emotion, Reason and the Human Brain*. London: Vintage.
- Deacon, T. 1997. *The Symbolic Species: The Co-Evolution of Language and the Brain*. New York, NY: W.W. Norton and Company.
- Donald, M. 1991. *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Harvard: Harvard University Press.
- Dunbar, R. 1996. *Grooming, Gossip and the Evolution of Language*. London: Faber and Faber.
- Evans, V. 2004a. *The Structure of Time: Language, Meaning and Cognition*. Amsterdam: John Benjamins.
- Evans, V. 2004b. How we conceptualise time. *Essays in Arts and Sciences XXXIII* (2): 13-44.
- Evans, V. To appear. *How Words Mean*. Oxford: Oxford University Press.
- Fillmore, C. 1982. Frame semantics. In *Linguistics in the Morning Calm*, 111-137. Seoul, South Korea: Hanshin Publishing Co.
- Fillmore, C. 1997. *Lectures on Deixis*. Stanford: CA: CSLI.
- Fillmore, C., Kay, P. and O'Connor, M.C. 1988. Regularity and idiomaticity in grammatical constructions: The case of *let alone*. *Language* 64: 501-38.
- Gallese, V. and Lakoff, G. 2005. The brain's concepts: The role of the sensory-motor system in reason and language. *Cognitive Neuropsychology* 22: 455-479.
- Garner, W.R. 1974. *The Processing of Information and Structure*. Hillsdale, NJ: Lawrence Erlbaum.

- Garner, W.R. 1978. Selective attention to attributes and to stimuli. *Journal of Experimental Psychology: General* 107(3): 287-308.
- Glenberg, A. and Kaschak, M. 2002. Grounding language in action. *Psychonomic Bulletin and Review* 9: 558-565.
- Goldberg, A. 1995. *Constructions. A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.
- Grady, J. 1997. *Foundations of Meaning: Primary Metaphors and Primary Scenes*. PhD Dissertation, University of California at Berkeley.
- Hurford, J. 2007. *Origins of Meaning*. Oxford: Oxford University Press.
- Jackendoff, R. 1992. *Languages of the Mind: Essays on Mental Representation*. Cambridge: MA. MIT Press.
- Kaschak, M. and Glenberg, A. 2000. Constructing meaning: The role of affordances and grammatical constructions in sentence comprehension. *Journal of Memory and Language* 43: 508-529.
- Lakoff, G. 1987. *Women, Fire and Dangerous Things: What Categories Reveal about the Mind*. Chicago: Chicago University Press.
- Langacker, R. 1987. *Foundations of Cognitive Grammar: Volume I Theoretical Prerequisites*. Stanford: Stanford University Press.
- Levinson, S.C. 1983. *Pragmatics*. Cambridge: Cambridge University Press.
- Mithen, S. 1996. *The Prehistory of the Mind: A Search for the Origins of Art, Religion and Science*. London: Orion Books.
- Prinz, J. 2002. *Furnishing the Mind: Concepts and their Perceptual Basis*. Cambridge, MA: MIT Press.
- Pulvermüller, F. 1999. Words in the brain's language. *Behavioral and Brain Sciences* 22: 253-336

- Pulvermüller, F. 2001. *The Neuroscience of Language: On Brain Circuits of Words and Serial Order*. Cambridge: Cambridge University Press.
- Schank, R. 1975. *Conceptual Information Processing*. New York: Elsevier.
- Schank, R. 1982. *Dynamic Memory: A Theory of Reminding and Learning in Computers and People*. Cambridge: Cambridge University Press.
- Schank, R. and Abelson, R. 1977. *Scripts, Plans, Goals, and Understanding*. Hillsdale. Hillsdale, NJ: Lawrence Erlbaum.
- Searle, J. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- Talmy, L. 2000. *Toward a Cognitive Semantics Volume 1*. Cambridge, MA: MIT Press.
- Tyler, A. and Evans, V. 2003. *The Semantics of English Prepositions: Spatial Scenes, Embodied Meaning and Cognition*. Cambridge: Cambridge University Press.
- Zwaan, R. 1999. Embodied cognition, perceptual symbols, and situation models. *Discourse Processes* 28: 81-88
- Zwaan, R. 2004. The immersed experiencer: toward an embodied theory of language comprehension. In *The Psychology of Learning and Motivation*, B.H. Ross (ed.), 35-62. New York, NY: Academic Press.