Lexical Knowledge, Conceptual Knowledge, and Generativity

Richard Carter

Various views exist concerning the relation between lexical and conceptual knowledge. Some argue for a gradation rather than a sharp distinction between the lexical and the conceptual [Haiman, 1980 ; Langacker, 1987]. Others accept the distinction, but disagree as to its nature. Some would include in lexical representations (LRs) information which others consider to be encyclopedic or conceptual, and different authors do this to differing degrees, thus the LRs of Jackendoff [1990] tend to be less encompassing than those of Pustejovsky [1995]. A related issue concerns the treatment of polysemy: should the lexicon specify multiple related senses for a given lexeme, or should the LR be such that multiple senses can be generated from a single representation, whenever possible, even at the cost of a more bloated set of LRs than more parsimonious theories would propose ?

In this paper I will address these two related issues, arguing first that a distinction between lexical and conceptual knowledge is motivated, and second that, while reduction of the number of senses that have to be explicitly represented for a lexeme is a desirable goal, minimizing the size of LRs is equally important. LRs can contain much less specification than some current views propose, since much of the work of a "generative lexicon" of the kind advocated in Pustejovsky [1995] is more appropriately handled at the conceptual level. The minimalist theory of lexicalization I advocate is more compatible with plausibles theories of the acquisition of lexical knowledge, and more explanatory in other ways.

1. The distinction between lexical and conceptual knowledge

I assume that knowledge is clearly *conceptual* if it is innately encoded as part of the human conceptual capacities, or is acquired independently of

language, it is clearly *lexical* if it concerns a particular lexeme and would not exist as knowledge if the lexeme was not known. These characterizations leave out many intermediate cases, but let's first consider cases that fit each description.

The innate basis for human conceptualization includes capacities for representing perceptual distinctions, colors, shapes, etc. These modular capacities exist prior to and independently of language. It is also often assumed that some more abstract representational capacities are innately specified, e. g. for spatial, temporal and causal relations and properties (cf. Bloom and Peterson [1997] for a variety of recent views about spatial representation). Most models of concept acquisition also assume that many concepts are acquired, some of these prior to and independently of language. Complex concepts like DOG, WASH, DOLL, may be acquired early on, perhaps as simplified protoconcepts involving linked prototypical representations created by innately specified perceptual-conceptual capacities.

It is is also frequently assumed in theories of lexical acquisition that the first LRs are acquired as simple connections between phonological forms (PFs) and protoconcepts, the "dog"-DOG theory ("dog" standing for the PF and DOG for the protoconcept). Such protoconcepts are clear examples of conceptual level (CL) knowledge, as characterized above. On many other issues there is less agreement, for instance, how much prelinguistic conceptualization there is, or how different protoconcepts are from "lexicalized concepts".

At the other extreme, consider next the case of two lexemes that are near or perfect synonyms, say "bunny" and "rabbit", or French "livre" and "bouquin". Presumably what we learn about such a pair of words is that they connect up to the same concept, say BOOK, and whatever differences there are. In these cases, there are differences of register: "bunny" is the child-language word, "rabbit" the unmarked word, for the concept LAPIN, and "bouquin" is the familiar word, "livre" the unmarked word, for the concept BOOK. Knowledge about register differences is clearly knowledge that we would not have if we did not know these words, thus a clear case of lexical knowledge, along with knowledge of phonological and syntactic characteristics. There is at least one more bit of lexical knowledge we have about these words, even assuming the "dog"-DOG theory: the knowledge of what concept each lexeme is connected to. This is acquired knowledge (the connection between PF and concept is arbitrary), and knowledge we would not have if we did not know the word. Thus, even a theoretician who asserts that there is no distinction between lexical and conceptual knowledge doesn't mean exactly that. What is meant, perhaps, is that the kinds of lexical knowledge I have described so far are all the lexical knowledge there is. That is, the only conceptually relevant knowledge we learn when we learn a lexeme is

what concept it connects to: LRs are always merely pointers to CL specifications, with, in some cases, register information.

But is this plausible? That depends crucially on what we believe the CL looks like. For instance, if we think that the CL has frame-like structures of the kind proposed in Fillmore [1985] and other literature on schemas, scripts, semantic fields, etc., then we are forced to admit that the lexical specifications of "buy" and "sell", for instance, are not merely pointers to CL structures, since they point to the same frame. The LRs must be different because "buy" and "sell" link to the commercial transaction frame in different ways, the subject of "buy" connecting to the node for the agent giving the money, etc., while for "sell" the subject connects to the other agent, etc. When lexemes connect differently to the same CL structure, the LR of each lexeme must indicate the differences. Here this gives LRs of the familiar predicate-argument structure sort. Furthermore, the specification that the subject argument of "cost" connects to the "goods" node, the object argument to the "price" node, etc., is knowledge that we would not have if we did not know the word "cost", knowledge relevant to understanding the word, hence lexical knowledge. Thus, if we accept that there are frame or field-size constructs at the CL, we are forced to admit non-minimal LRs for such words.

There are other plausible cases of non-minimal LRs, besides those involving situation frames. Consider the word "horse". Suppose we assume that the LR of this word is a bare distinguisher [Katz, Fodor, 1963], that is, a pointer to the CL structure HORSE. A near synonym like "nag" (a derogatory term for horse) should point to HORSE as well, even if the meanings are not exactly the same, for all of the CL information relevant to the understanding of "horse" is relevant to the understanding of "nag". Our specification of the meaning difference should not mask the essential identity at the CL. Similar considerations should lead us to the same position for "mare", "colt", "pinto", "steed", etc. If we wanted at all costs to maintain that an LR has nothing but a distinguisher pointing to a concept, that would force us to posit HORSE, MARE, PONEY, etc. as distinct concepts. We would then have to organize the CL so that these minimally different notions could inherit almost all of the same body of knowledge, about prototypical shape, eating habits, etc.

Is MARE in fact a plausible long term memory concept? If lexical knowledge is knowledge we would not have if we did not know the lexeme, then the knowledge that "mare" designates female horses looks like lexical rather than CL knowledge. Consider "kangaroo". Like "horse", "kangaroo" applies to both males and females, and as with horses, we know that there are both female and male kangaroos. But English has no word for female kangaroos, so if we want to refer to them, we have to use a complex expression like "female kangaroo". Suppose it did have such a word, say "froo". Then what we would know about "froo"

is that it is referentially restricted to female kangaroos. Not having such a word, we do not have that particular bit of knowledge. On the other hand, we may have specific knowledge about female kangaroos (e. g. that they have pouches), as we may about female horses. The only clear knowledge we have about "mare" that is not parallel to the knowledge we have about female kangaroos is the knowledge that "mare" denotes female horses. There seems to be no more reason to posit a concept MARE than there is to posit a concept FROO for female kangaroos: the existing concepts KANGAROO and FEMALE are enough to handle any FROO thoughts, and similarly for MARE thoughts. What is specific about "mare" is that it triggers the complex CL representation FEMALE HORSE. Note that, even if we think that there is a semantic difference between "mare" and "female horse", that would not of itself justify positing the CL entity MARE: if the difference can be explained by general principles, then the distinction can be produced generatively at the CL, and does not have to be encoded as a concept in long term memory. Finally, note that a "mare" is necessarily female and a horse. Thus the alternative to maintaining the "all LRs are simple distinguishers" theory at the cost of positing implausible CL constructs, is to accept that the LR of "mare" contains not only a pointer to HORSE, but also a specification that the denotata must be female.

We see again that the question of what LR's look like cannot be treated in isolation from the question of what the CL is like: if CL representations can be frames, then some LRs will resemble predicate-argument structures, if multiple lexemes with distinct referential constraints, like "horse", "nag", "mare", can point to the same CL concept, then LRs have to contain referential specifications. In short, the theory of possible LRs cannot be developed independently of the theory of possible concepts, although our argument indicates that these theories are distinct, that LRs are not merely chunks of conceptual representation.

Such considerations place a lower bound on LRs: unless we assume that concepts are always in a one-to-one relation to lexemes, LRs are not always bare pointers to CL structures. Clearly, the less CL structures resemble lexemes, the more need there will be for non-minimal specifications of LRs. But what about the upper bound? How complex can LRs be? I will now argue that LRs should not be much more complex than the above examples suggest, that is, that LRs are basically pointers to CL representational structures, with, in some cases, additional constraints on how they are linked to the CL. The function of LRs is to activate CL structures, and to configure them in certain limited ways. However, the representational services that provide understanding are furnished at the CL, the role of LRs being simply to activate that level appropriately. This view of LRs as an interface to conceptualization is reminiscent of that of Jackendoff [1997] but, I shall argue, the LRs it calls for are rather different from those developed in, for example, Jackendoff [1990].

If the only function of LRs is to activate CL structures, then clearly the nature of LRs depends on the nature of the CL. Many have argued for views of the CL on which it is not merely a box of lexeme-sized concepts, but has larger-scale structuring of one kind or another, and I also believe this to be the case. I cannot present and defend here in detail the modular view of CL structure I find most plausible, but a brief outline may make it clear why a minimalist view of LRs and modular view of conceptual structure is plausible.

The Structure of the Conceptual Level

One version of the view that the CL has larger-scale structures is presented in Lakoff [1987] where it is argued that lexemes are associated with Idealized Cognitive Models (ICM), structured in a variety of ways. These range from ICMs that represent the kind of knowledge brought to bear in the understanding of calendric terms like "month", "day", "January", etc., where one ICM is activated by many distinct lexemes, to a case like "mother", in which many distinct ICMs can be triggered by a single word, "mother". I will not attempt a full critique of Lakoff's work, but offer a few remarks aimed at pointing up differences between his views and a different one, that may help make the following sections more plausible.

For calendric terms, it is clear that the body of knowledge we bring to bear in using these terms is highly structured, in fact generative. The notion of generativity is crucial here, and, I believe, in general: that is, the CL is organized in a highly articulated modular architecture of numerous generative systems, and the function of LRs is to appropriately activate these systems.

It is eminently plausible that our calendric knowledge constitutes a large complex generative system, composed of several sub-systems, which enables us to produce an open-ended set of calendric mental representations. These are used, for instance, in solving problems like "if John left Paris at noon on January 1, 1996, and had arrived exactly 24 hours before, when did he arrive?". The set of problems we can solve is open-ended, and so is the set of partial calendric-chronometric representations we can construct to solve them, although of course restricted by memory limitations. A thorough and cognitively plausible formalization of this knowledge as a generative system remains to be developed (cf. [Zablit, 1991] for a brief discussion).

I believe that generativity is the general case. Any "concept", or CL structure, is profitably regarded as an analytic-representational generative capacity, whose function, whether activated by language, perception, other concepts, or memory, is to produce an open-ended set of representations. For instance, the commercial transaction "frame" can be regarded as a

generative system for producing an infinite set of representations of buying and selling situations, etc., rather than a static slot-filler structure. Furthermore, what is very plausible about the case of calendric knowledge is that it is a representational capacity that we could not use without the intervention of linguistic representations, and could not acquire without explicit instruction crucially involving language. That is, it is a CL capacity that depends on the possession of linguistic capacities.

It is clear that the knowledge recruited in this capacity is not exclusively lexical [Leech, 1969]. For instance: we know how to apply day and month names serially, know, for instance, that Thursday follows Wednesday and precedes Friday. We know that the system of nomenclature is iterative— after Saturday comes Sunday, then Monday, and so on forever; in fact we know this even though we may not be sure whether the week "begins" with Sunday or Monday. We know how days relate to weeks, weeks to months, months to years, how many hours there are in a day, how many days in January, etc. None of this knowledge is of the kind anyone would include in a LR. But, to use the capacity, we must connect lexemes appropriately to the knowledge, for instance, connect "day", "holiday", and "Monday" to the knowledge about days, etc. Some of the terms ("leap-year", "hour", "noon") have more explicit characterizations than others ("weekend", "holiday", "twilight"). Furthermore, details of lexicalization vary from one language to another. For instance, corresponding to "day" Swedish has both "dag" (day), and "dygn", which refers more explicitly to the twenty-four hour period, French has the noun "lendemain" meaning "following day", contrasting with the adverb "demain", meaning "tomorrow". These are examples of lexical knowledge, knowledge about words: an English speaker has no knowledge about any English word corresponding to the French speaker's knowledge of what "lendemain" means, although he knows what the expression "the following day" means. Similar points can be made about commercial transaction knowledge, and many other language-dependent systems used in thinking.

It thus seems necessary to distinguish between two types of "language-dependent" knowledge, i. e., knowledge one would not have if one did not have the word. First, there is CL knowledge like calendric knowledge that, arguably, we could not acquire and represent if we did not have language, this is CL, not linguistic, knowledge, best characterizable as a language-dependent generative system. Second, there is lexical knowledge. For example, both English and French speakers understand the concept of one day following another, but only French speakers have the lexical knowledge of what "lendemain" means, while English speakers express this conceptualization with expressions like "the following day", "the next day" (or the archaic "the morrow").

I would argue that not only calendric knowledge and explicitly taught systems (geographical, arithmetic, etc.) are cases of language-dependent

systematic CL knowledge, but much other CL knowledge is of this kind. Be that as it may, for such knowledge, it is necessary to distinguish between lexical and CL knowledge, if we agree that the CL knowledge is structured. Whether the structures are Frames, Schemas, ICMs, or generative systems, there is a need for LRs that are more than unarticulated pointers to the CL system: thus "lendemain", like "mare", must trigger a CL configuration, the same one activated by "day", and must impose the additional requirement that the day denoted immediately follows some other contextually specified day. For example, "Il est arrivé jeudi et il est parti le lendemain" means "he arrived Thursday and he left the next day" (i. e., Friday).

A second, rather different, example of an ICM is the complex of ICMs associated, in Lakoff [1987] with the lexeme "mother". I will mention here only one of the many points that can be raised about this case: Lakoff mentions five ICMs that can be relevant to the interpretation of "mother" in different contexts, these ICMs forming a "radial structure". They include "the birth model" ("the person who gives birth is the *mother*"), "the nurturance model" ("the female adult who nurtures and raises a child is the *mother* of that child") etc. It is striking that all five ICMs share certain specifications, for instance they all involve relations, between x and y, both human, in which x is a female individual and the sex and number of y is indifferent, and in which there is a causal relation of some kind going from x to y (x may be the individual who provided the egg that ultimately combined with a sperm cell to yield y, or the individual who raised and cared for y, or the person who gave birth to y, etc.). Lakoff argues for the existence of different ICMs because there are different sets of assumptions about the relation involved. As he would no doubt agree, the different knowledge systems evoked by these ICMs are themselves complex generative systems that also figure in the interpretation of other lexemes (e. g. "child", "parent", "mommy", "grandmother", "birth", etc.). What is left unexplained from the structure he proposes is why all these ICMs should have the similarities just noted.

I suggest that they do because there is a common core of "meaning" to the word "mother", a specification associated with a unitary CL structure closely related to the word, which includes the information that a mother is female, bears some causal relation to a typically younger person, etc. This structure in turn must relate to other structures, perhaps similar to Lakoff's ICMS, that are also activated by other concepts. A theory of the relations involved in understanding uses of these words must, as Lakoff proposes, involve the action of rich CL systems, whose nature and functioning remain to be worked out in detail.

I think that this example, like the previous ones considered, rather than supporting the claim that there is no distinction between lexical and CL knowledge, argues in favor of the distinction. These examples support a model of LR in which lexical specifications rarely if ever constitute

necessary and sufficient conditions. What is striking about these examples is that there are a few necessary conditions, that greatly underspecify our intuitive notion of the "meaning" of the words. For instance, if all that is necessarily true about a mother is that she is female and causally related to another individual, that does not distinguish "mother" from "aunt" or "murderess". As is well known, lexemes frequently trigger judgments of necessity, but rarely allow definition, i. e., rarely entail an expression that entails them. The proper model of the relation between lexemes and conceptualization must *explain* both why words don't have definitions and why they sometimes have necessary content. I know of no model that does this, but I will suggest an approach that seems promising.

2. Towards a minimalist theory of lexical representation

The language learner must learn arbitrary connections between PFs and the CL. In the case where a CL concept (or concept-sized representational capacity) already exists, the infant might simply connect the PF to the concept, "dog" to DOG. That is the easiest thing to do, so let's assume that the learner will try to do this whenever possible. It will be possible if there already exists a plausible candidate concept, otherwise not. Let us now consider two simple cases where there is no plausible candidate, first, the case where there is no candidate at all, and second, the case where the CL representation is not of the right size, for instance where the CL representational capacity is a schema-like structure.

Presumably there will often be cases of the first kind. For instance, suppose the child's first exposure to the word "kangaroo" involves hearing the word and seeing an instance at the zoo or in a picture book. The child must simultaneously create the CL representation and associate it with the PF. Let us assume, very simplistically, that something like this can happen: the child creates a CL structure of the kind he might have created prelinguistically (as he did for DOG) but with linguistic input. This scenario seems to pose no new problem, unlike putative cases of acquisition in which a concept could not be acquired without linguistic support (e. g. calendric terms, or theory-laden words like "buy", "money", "phlogiston", "lexicalization", "sin", etc.).

Consider next the example in which the child must connect a PF with a CL structure that exists, but is not of the right "size". This must be a frequent case. For instance, with color terms, undoubtedly a rich prelinguistic representational capacity exists, providing various candidates, and the problem for the child is to pick the right one. Current proposals about this case assume a hierarchy of candidates (focal colors first) and perhaps some field-type principle that determines the extension of a term as a function of the existence of other lexicalizations (to explain, for

instance, why the child who learns "rouge" in French, which distinguishes red from orange, doesn't extend "rouge" to orange hues, while a child who learns a language with no separate word for orange does, on the assumption that both learners associate the word learned with focal red).

This kind of lexical difference between languages is extremely common. For instance, where English has the single word "chair" whose extension covers a certain range of artifacts, the question of whether they have arms or not being irrelevant, French has two words, "chaise" and "fauteuil", which divide up roughly the same extension in such a way that whether the thing has arms or not is relevant, "chaise" being most typical for armless chairs, "fauteuil" for chairs with arms. The facts are complex : for instance, theater seats, or car seats, are "fauteuils", but not "chairs", a "wheelchair" is a "fauteuil roulant", and a "chaise électrique" typically has arms but is not a "fauteuil" (perhaps because "fauteuil" has inappropriate connotations of comfort, or because the term is a loan translation).

The problem for a theory of acquisition is to specify a plausible model of the innate basis and of prelinguistic conceptualization, and show how the learner of French or English acquires the knowledge : is there, for instance, a common protoconcept that the learner of French and English adapts differently? That is conceivable, since there seem to be no relevant non-linguistic differences between furniture in French and English-speaking environments —each language just seems to cut up the environment differently. If there is a protoconcept, CHAIR*, is it indeterminate with respect to the differences between the languages? Or is CHAIR* more like "chair", or "chaise", or "fauteuil"? Or does the theory of prelinguistic conceptualization lead the learner to make a distinction similar to that distinguishing "chaise" from "fauteuil", creating the protoconcepts CHAISE* and FAUTEUIL*; so that the learner of French has little to learn except that "chaise" connects to CHAISE* and "fauteuil" to FAUTEUIL*, while the English learner must learn that "chair" connects to both? Even with two languages as closely related historically and culturally as French and English there seem to be hundreds, if not thousands, of such cases (e. g. F. "savoir", "connaître" vs. E. "know"; F. "bus", "car", vs. E. "bus", F. "à", vs. E. "to" "at", etc.). Such cases suggest that prelinguistic conceptualization leaves a fair amount of leeway for differences in lexicalization, which in turn suggests that the child must work harder than the simplest model would claim, that he must develop and differentiate lexemes and conceptualizations together, rather than representing LRs as simple pointers to existing concepts.

Thus the simplest plausible model of lexical acquisition may be one in which the learner adopts simple pointer LRs whenever possible, and, when this is not possible, gives LRs the minimal specification that will connect the lexeme appropriately to the CL. That is, the learner does as little as possible. The empirical problem is, what is "as little as possible"?

If the learner of French has no conceptual distinction corresponding to the difference between "chaise" and "fauteuil", then he has no option, if he is to use the words appropriately, but to create the conceptual distinction. If, as is plausible, he has already noticed that some seats have arms, some not, some are soft, some hard, some are made of wood, some of metal, some fold up, some do not, etc., he must decide which of these distinctions is relevant to these two words, and encode it in long term memory in connection with these words. A relevant observation is that French speakers know that some of these distinctions are irrelevant, while others are relevant. For instance, possession vs. absence of arms is highly relevant, and comfort, hence softness, is also relevant, while curved vs. straight legs, color, or shape of back, are of no relevance. As usual, while speakers have clear judgments about prototypical examples, they are less sure of more marginal ones. One interesting observation is that a folding garden chair, made of plastic, with quite comfortable arms, is an inappropriate fauteuil, but a fairly respectable "chaise", although not typical, while for a stuffed or padded salon chair, with no arms, neither "chaise" nor "fauteuil", but some other word, such as "chauffeuse", is preferred. That is, arms seem to be required for "fauteuil", but are not sufficient. On the other hand, if a seat has neither arms nor a back it could not be called either "chaise" or "fauteuil", but "tabouret" (approximately, "stool") or "banc" (approximately, "bench"). There are comparable facts for the distinction in English between "bench" "couch", and "sofa". For instance, none of these absolutely must have a back, but "couch" and "sofa" require a certain width, enough to allow several average-width humans to sit on them. Having a back seems closer to a necessary requirement for "sofa" than for "couch", but if something is made out of wood or metal, with no padding, it could be called a "bench", but hardly a "couch" or "sofa", even if it has a back. Again, the French terms "banc", "canapé" are similar to but not exactly equivalent to any English words.

Such observations suggest that, if it is hard for the linguist to come up with a reasonable representation of the "meaning" of such terms, distinguishing between normal and marginal cases, it is perhaps equally hard for the language learner, who may decide on only partial or provisional descriptions. That is, he may store representations of the most commonly encountered instances as visual schemata, perhaps assigning one or two features to the LRs of the words themselves, and leave it open what the status of observed traits (e. g. comfort or padding, folding) is. A theory of lexical acquisition must allow for incomplete acquisition, and for some degree of sophistication in the acquirer, since there are cases for which we do not know, and know that we do not know, what the information that determines the extension of a word includes ("elm" is a standard example).

Theories can be distinguished according to whether they say that there is a distinction between lexical and CL knowledge; and also, for those that

admit the distinction, according to whether the learner may know, i. e., represent, that he does not know the LR, or knows it only partially. If the theory says he never even looks for a LR, because there cannot be any, then it can deny to the learner this kind of sophisticated epistemological attitude. In the present model, the lexical acquisition device looks for highly consistent information about how a word is used, and when it decides it has found some, it incorporates it in the LR. That is why all lexical knowledge must be "true by definition".

At the other extreme, we might expect conceptual knowledge which is hard-wired into the device, i. e., innately specified, to also be treated as certain, in some sense. That is, CL hard-wiring should have a different status from the CL knowledge that is acquired through experience, which should often involve merely relative certainty, since our concept acquisition mechanisms are responsable for this knowledge, and should, (at least if we were optimally designed) be ready to revise it. We do seem to be capable of revising acquired knowledge systems, e. g. what we believe about horses, more easily than what we believe about space and time.

To sum up, two distinct sources of judgments of necessity, or at least our most automatic judgments, should be (1) lexical knowledge, for *trivial necessity,* knowledge true merely by definition ; and (2) our innately specified modules, hard-wired in the way our visual system is, for automaticity. Thus our judgments about the nature of space and time, for instance, seem to have a similar status to our judgments about perceptual illusions. We may judge, by reason, that the lines in Muller-Lyer figures are the same size, but we can't help seeing them as different. Similarly, perhaps, we may persuade ourselves by verbal reasoning that space is curved, and bounded, or that time had a beginning, but we can't help conceiving space as Euclidean, or time as unbounded.

A minimalist theory of lexical acquisition, hence of lexical representation, such as that suggested here, may be of explanatory value. Not only would it help explain why lexical content yields judgements of necessity: the less information a LR contains, the less reason for it to contain non-essential information. It also allows us to predict that LRs will depart from the minimal case (where the LR is merely a pointer to a concept) only under restricted conditions. For instance, a word like "mare" should only exist if certain other words already exist. Suppose that the learner hears the word "mare" applied to a female horse, and has the protoconcept HORSE*, but knows no word for it. He will automatically, on a minimalist theory, assume that "mare" means HORSE*, since that is the minimal hypothesis. The only situation in which he could represent a word as meaning female horse would be one where he already knows at least one another word, either one meaning male horse, or one meaning horse. The general prediction is that words with complex LRs can be learned, hence can exist in someone's mental lexicon, only if either other

words of simpler meaning are already known, or other words of no more complex meaning are simultaneously acquired, and the combined set of lexemes covers the CL structure. This suggests that the kind of hierarchy of lexicalization proposed by Berlin and Kay [1969] should be found in all conceptual domains.

To evaluate this proposal about how lexicalizations are structured, we need a developed theory of the CL and of concept acquisition, which I will not attempt here. As examples of what the proposal implies, let us consider briefly again two cases already discussed, commercial transaction lexemes and horse-denoting nouns. The constraint says that the only way a lexeme can come to have a LR that is more than a simple pointer to a conceptual site is for it to be learned in contrast to other lexemes pointing to the same site. The learner will include further specifications in the LR when he has to distinguish the lexeme from other lexemes connecting to the same site. If the site is the commercial transaction frame, then frame-activating words like "buy", "sell", "price", "cost", must be distinguished from each other, since they must activate the site differently, so their LRs must be non-minimal. Likewise, if "horse", "mare", "colt", "filly", all point to the site HORSE, then their LRs must be distinguished, since they must activate it differently. The LR of "horse" might simply point to HORSE, while the LR of "colt" might also include the specification YOUNG, etc.

If it is correct that the CL consists to a great extent of larger-than-lexeme structures such as frames, fields, ICMs, as argued by many (e.g., field theorists, Lehrer, Lehrer and Kittay, Fillmore, Lakoff, Langacker), then many LRs will have to consist of more than a simple pointer. The prediction is that all such lexemes will have "definitional"[1] content, and this seems to be the case. Note that "horse" has no such content, while "mare", "stallion", "colt" do, but that the definitional content does not represent all we know about typical mares, colts, etc. (for instance, mares are adults, colts are young and frisky).

[1] I use the term "definitional" here in a special sense, for LR specifications, with no implication that the specifications constitute a definition, avoiding the alternative "conventional", which I save for a more general sense. For instance, the knowledge that February follows January is conventional, but it would be odd to say it is specified in the LR of either word. This is another kind of "necessary" knowledge in addition to those mentioned below.

3. Where's the information ?

If such a minimalist view of lexicalization is on the right track, some current proposals about LR cannot be correct. I will examine two of these here, first discussing an analysis by Ray Jackendoff of the semantics of "climb", and then consider the type of analysis proposed in Pustejovsky, 1995, in the "generative lexicon" framework. I will argue that in each case a more explanatory account can be achieved if some of the material included in the LR is handled not in the lexicon, but at the conceptual level.

In the analysis in Jackendoff [1985, 1990] of the English verb "climb" it is claimed that the LR includes specifications, among others, of upward movement and of "clambering", and that while it is preferable for both aspects of the meaning to be present in the situation denoted, neither aspect is obligatory, rather what is obligatory is that one of the two be present. I believe this description is not quite accurate, and conflates distinct phenomena. There are perhaps three relevant cases, illustrated in (**1**)-(**3**) below:

(**1**) *John climbed.*

(**2**) *John climbed the hill.*

(**3**) *John climbed down.*

In (**1**) "climb" is used absolutely, in (**2**) transitively, and in (**3**), although used intransitive as in (**1**), "climb" is construed with a particle. The facts are these: (a) When used absolutely, as in (1) "climb" strongly implies upward movement (although perhaps that implication can be overridden by context), and also a hard-to-pin-down causal notion which Jackendoff calls "clambering", and which I will label as "effort", and which seemingly is always present (can never be overridden by context). This causal effect is observable in other absolute uses:

(**4**) *The plane climbed rapidly.*

(**5**) *The balloon climbed for ten minutes.*

(**6**) *The snake climbed.*

(**7**) *The temperature climbed.*

It is revealing to compare these uses with "rise":

(**8**) *The plane rose rapidly.*

(**9**) *The balloon rose for ten minutes.*

(**10**) *The snake rose.*

(**11**) *The temperature rose.*

Connotations and relative naturalness vary with choice of subject: (**4**) suggests extra effort, speed or height, while (**8**) suggests effortlessness, or neutrality of description. (**5**) may seem odd, but could also be used to suggest rapid ascension, perhaps as the balloon is pushed by the wind. (**6**) could describe a snake's upward movement as it crawls up a branch, while (**10**) can evoke the picture of a snake floating upward, or a clinically neutral description; (**11**) is a colorless description, but (**7**) is marked, suggesting perhaps rapid or unusual rise.

What (**4**)-(**7**) have in common with (**1**) is that they entail upward movement.

Consider, now, (**12**):

(**12**) *John climbed for an hour, but he still did not reach the bottom.*

This seems special, even though, as Jackendoff argues, it may be possible to use it for downward movement involving "clambering", but only in a context where it is clear that the movement is downward. I don't think that (**12**) all alone provides such a context, since, "the bottom" could refer to the bottom of a ladder or chimney up to which John climbed, and unless there is extralinguistic context indicating downward movement, (**12**) would be more natural with "downward".

These observations indicate that (i) upward movement is a prominent feature of absolute "climb", and (ii) there is always some other connotation with absolute "climb" that distinguishes it from "rise". Note that, if intransitive "climb" could be used absolutely with just one of the meaning components (UP or EFFORT) without the other, we would expect it, in some cases, to be identical in meaning to "rise", but that does not seem possible: sentences with "climb" differ in meaning from sentences with "rise".

Consider now transitive "climb", which can be used with NP objects like "hill" or "stairs" that designate the physical site of the climbing, or with NPs indicating the distance climbed, as in (**13**):

(**13**) *John climbed the hundred meters in ten minutes.*

In both cases, there is an entailment of upward movement, and a connotation of effort, plus an entailment of totality: in (**13**), the upward movement covers the whole distance, in (**2**), the whole vertical length of the hill. This is the reason Jackendoff includes a specification TO TOP in the LR of "climb". However this information seems predictable, hence unnecessary: there apparently could not be a transitive verb that was identical in meaning to "climb" except that in (**2**) or (**13**) it could describe movement over part of the distance. Furthermore, this seems to be a general characteristic of transitive verbs of this type as can be seen with transitive uses of "walk", "hop", "swim":

(**14**) *He swam the pond.*

(**15**) *He walked the distance.*

(**16**) *She hopped the fence.*

and is presumably related to the reason why, in (**17**) but not (**18**) there is an entailment that the wagon is filled with hay:

(**17**) *She loaded the wagon with the hay.*

(18) *She loaded the hay into the wagon.*

If we simply specify the LR of transitive "climb" as requiring upward movement, and it is predictable that the action of the verb concerns the relevant totality, it will follow that "climb the hill" entails upward movement from the bottom to the top of the hill (incidentally, Jackendoff's LR does not actually stipulate that the movement must start at the bottom, as seems necessary).

Consider now the third case:

(3) *John climbed down.*

In this use, "climb" does not entail upward movement, in fact (3) specifies downward movement. That is the most obvious reason why it might seem justified to specify UP in the LR of "climb" as an overrideable preference, rather than a necessity, in spite of the a priori desirability of not having to stipulate the difference between non-overrideable and overrideable specifications in LRs.

But this case is different from the others, since the very possibility of combining "climb" with a constituent like "down" is a marked phenomenon of English [Talmy, 1972, 1985 ; Carter, 1980 ; Jackendoff, 1990]. That is, in certain languages, including English, it is possible to use a verb that expresses manner of movement, like "walk", "swim", or "climb", with spatial PPs such as "under the bridge", and particles such as "down", "home" even when these are not lexically specified for movement. This is not possible in other languages. For example, in French:

(19) *Jean a marché sous le pont.*

can be used to describe an incident of walking in the space under the bridge, but not a displacement from a point not under the bridge to a point under the bridge. That is, "sous le pont" cannot express the "goal" or "path" of movement.

I believe the correct description of this difference should not say that French excludes in general the combination Verb of manner of movement plus path PP, since in fact such sentences as the following are perfectly normal:

(19) *Jean a marché jusqu'au pont.*

(20) *Jean a nagé jusque sous le pont.*

(21) *Jean a marché de Paris à Versailles.*

These examples show it is possible to combine verbs like "marcher" with a PP as long as the PP expresses a path, as is the case with "jusqu'au

pont" or "de Paris à Versailles". What is not available in French is the marked possibility of describing movement along a path with a non-path PP such as "under the bridge" and a manner verb like "walk", "climb", or even "float". I suggest that this extra flexibility of English can explain why it is possible for the semantics of "down" to override the UP in the LR of "climb".

Note that French "grimper", the nearest equivalent of "climb", can be used only intransitively, and always entails movement upwards, so that "grimper jusqu'en bas" can only describe movement upward to the bottom of something not mentioned, say a ladder (In fact the Petit Robert's first sense of "grimper" speaks of moving upward using the hands). For transitive "climb" the closest French equivalent is "escalader", which, as we expect, entails movement from the bottom to the top of the object.

Prelinguistically, future learners of French and English should have the same concept and knowledge, for instance that entities can move in different directions, and that upward movement requires more effort than other directions. If one group has a version of the CLIMB concept then so should the other. But on Jackendoff's analysis of "climb", upward movement is not an obligatory part of the meaning of the word, so why for French "grimper" is UP obligatory? If, that is because for the prelinguistic concept upward movement is obligatory, then how and why does the child who learns English "climb" weaken that specification?

On the other hand, it is not clear why there should be any prelinguistic conceptualization for CLIMB at all: why should a conceptualizer have singled out the case of moving upward with effort for specific conceptualization, as opposed to moving downward without effort, or horizontally, or for all sorts of combinations that neither language lexicalizes? Plausibly, the prelinguistic child can represent movement in *all directions, using hands, legs, crawling, and can generate* representations combining these features, and has no need to encode any particular combination as a concept in long term memory. And even if the child has noticed that movement upward requires extra effort, it is not clear why that should motivate creation of a concept CLIMB, instead of a general specification that ease of movement varies inversely with direction.

If there is no prelinguistic concept CLIMB, then the child learning "climb" or "grimper" must create the concept. On our assumption in such a case the child might simply create a LR with the specifications UP and EFFORT (to distinguish "climb" from "rise") and these specifications will automatically be definitional. That explains why "grimper" entails upward movement. An alternative model has to explain how the child figures out that "grimper" entails upward movement, while for "climb" UP is only a preference. On the suggested analysis, both intransitive "climb" and "grimper" receive the same LR, with UP and EFFORT, and the possibility of "climb down", or of an elliptical interpretation of absolute "climb" in

certain rich contexts is a consequence of the special interpretative possibility for manner verbs. This approach to explaining the facts implies that lexicalization leads to necessity, but if it does here, then plausibly it does so in general.

In fact, on a priori grounds, it seems reasonable that lexicalization should induce a presumption of definitionality: the existence of the lexeme implies that speakers are interested in that particular type of activity, although the learner may have to figure out just what that type of activity is. Understanding and dealing with the world presumably interests the child more than how people talk about it, so it seems natural that it would prefer the assumption that the relation between the lexicalization and what it names is "true by definition". The problem is that often it is not clear just what the scope of the word is. This might suggest a cautious strategy in which the learner tries to take care of the clearest aspects of a word's use, taking them to be definitional content, leaving other aspects for further consideration. That seems to be the usual pattern: there are thousands of words that have a minimal set of clear entailments, and are associated with a larger, indeterminate, and evolving set of image schemas, typicality conditions etc. that play an important but non-determinative role in their use and understanding.

On this view, then, a major source of the certainty judgments connected with words is decisions that some aspects of word use don't have to be worried about. Such decisions could well be revisable over time (cf. [Carey, 1985] for the changing status of surface features of animal kinds, apparently from definitional to incidental).

4. Where's the generativity ?

Among the current models that accept the validity of the distinction between lexical and CL knowledge, the generative lexicon (GL) model most massively violates the stricture that all lexical specifications be definitional (non-overrideable). For instance, in Pustejovsky [1995] the LR for "book" specifies that books are physical objects that contain information, are read, and written. None of these characteristics are definitional, although they are easily accessible aspects of our knowledge about books, and relevant to the interpretation of the word "book". On the other hand, other easily available knowledge is not included in the LR, e. g. that books are usually printed nowadays, typically have covers, are read from front to back, have typical sizes, shapes, etc. (few are pretzel-shaped, ten meters high, made of uranium, written in orange ink on red paper). As Katz and Fodor [1963] pointed out, anything we know about the denotation of a word can be relevant to its interpretation; they distinguished between a subset of "dictionary" information that was to be

included in the LR and all the rest, the "encyclopedic" information. The difficulty of formulating criteria for this distinction has led many to reject it altogether. What is novel about the GL approach is that it accepts the dictionary-encyclopedia distinction, but it draws the line not in terms of the distinction between necessary and non-necessary information, but in the light of the requirements of a set of generative mechanisms supposed to operate over LRs to produce "senses". The approach raises interesting and difficult problems for an account of how such LRs are acquired.

One such problem arises from the fact that some of the specifications correspond to definitional content, while others do not. For instance, it is presumably definitional that knives are for cutting, and this is expressed by the TELIC specification **cut(e,x,y)**. On the other hand, while books often contain text, and are meant to be read, that is not always the case, as Pustejovsky points out: a book can contain only pictures, or numbers, or blank pages, and still be a book, as are atlases, notebooks, musical scores, etc. Thus the TELIC specification **read(e,x,y)**, has a different status than that for "knife". Speakers know this, but the knowledge is not represented in the example LRs, although of course it could be, non-explanatorily.

A similar case is that of "window". The GL account of how we understand "open the window" or "enter through the window" proposes a LR for "window" parallel to that for "door": both specifying that there is an aperture. The LR for "door", [Pustejovsky, 1995, p. 223] specifies both physobj.aperture and TELIC "walk_through(e,z,y)". The TELIC specification for "window" will hopefully not be walk (or jump) through, but a more reasonable candidate like "see through". Presumably by definition windows are for seeing through, and it is not necessary that they have openable apertures, as in fact many do not. We know that they are sometimes openable, and that therefore a possible interpretation of "the bird flew through the window" is that it came through the aperture left by the open window, and this is why the LR of "window" will specify that it is an **aperture** (p. 93). These are typical examples of how the GL approach adopts LRs with specifications that are not always true.

An obvious alternative to the GL approach is one in which interpretations are constructed on the basis of CL specifications. For instance, in understanding "jump through the window" we use our knowledge that windows are often openable in combination with other knowledge to obtain as a possible interpretation that there was movement through an aperture. On this view, LRs are just one of the factors that enter into the construction of an interpretation, and a less important one than in the GL model.

One consideration motivating the GL approach is the desire to avoid the extensive polysemy of some other approaches. The idea is to have a single, albeit hefty, LR, and to generate "senses", rather than simply listing them in LRs. But an alternative would be to have unitary and more modest

LRs, and do the generating at the CL. This would imply abandonment of the claim that there is a purely compositional semantic level. On the minimalist view of lexicalization sketched above, it would make little sense to have compositional processes operating at a purely lexical level, for there would be little for them to do. The natural way to view LRs would be as activators of and constrainers of CL representational activity. I will now argue that this alternative will have more explanatory power that the GL approach.

Consider the sentence:

(22) *John has finished the book.*

One possible, no doubt frequent, interpretation is that John has finished reading the book. The set of potential interpretations, however, is infinite. For instance, if we know John is a printer, we may understand the activity to be printing. In fact, that might well be the automatically produced interpretation of:

(23) *The printer has finished the book.*

Furthermore, by replacing "printer" by "illustrator", "book-burner", "termite", "book-eater", etc. we can easily trigger other interpretations. One claim of GL theory is that a small set of the potential interpretations, those encoded as the Qualia of a noun like "book", has a special status in the interpretive process.

It is not clear what could motivate the learner to put these non-definitional specifications in the LR, nor is it clear what justifies the claim that they have a special status, other than that of being activities that frequently come to mind when we think of books. Their special status cannot be that they force the interpretation, since they do not. A crucial observation militating against this claim is simply that in such sentences *contextual knowledge always overrides all lexical or other CL factors.* Thus, if we know that John is a printer, but that what is being discussed is a book-eating contest, the automatic interpretation of "John finished the book" or "the printer finished the book" will be *finish eating,* no matter how unusual the activity of book-eating is. The only conceivable special status for such LR specifications would be to automatically provide plausible candidates for the activity. This would not be a legitimate argument for anyone purporting to provide a theory of semantic *competence,* but it could be an argument for a performance model: encoding the most likely candidates in the LR might facilitate interpretive processing. However I know of no evidence (such as, e. g. backtracking effects) that would justify such a claim, which seems unlikely, since our CL capacities for constructing interpretations are quite powerful, able to deal with ellipsis, figurative language, etc. In fact, it is a commonplace that

the inability of current computational linguistics to model the non-linguistic aspects of the interpretive process is the major explanation for the poor performance of natural language understanding programs. It is not clear that our conceptual level capacities need the limited kind of help a GL could offer.

Be that as it may, the fact that LR specifications can only offer candidates for the missing activity, and that context-specific knowledge always overrides such candidates poses a problem for the GL model: why does context override this kind of LR specification, and not all LR specifications? There are, in fact, some LR specifications that context cannot override, and these seem to be the ones that a minimalist model would say are the only lexical specifications. Consider "finish the book" again: as Pustejovsky points out, "finish" imposes a specific constraint on the type of activity needed. Roughly, it must be an "accomplishment" (see Pustejovsky [1995] for a more precise description). This constraint is clearly inviolable. Consider, for instance:

(23) *John finished the chair.*

Presumably the TELIC specification of "chair" will indicate that a chair is for sitting, but sitting is not an accomplishment, and so this cannot have the interpretation "John finished sitting in the chair". In fact no priviledged activity automatically comes to mind here, but we can easily imagine many: painting, sanding, assembling, photographing. The claim of the minimalist is that all LR specifications are inviolable. This offers a straightforward explanation of why in "finish the book" *reading* is not a necessary part of the final interpretation: the knowledge that books are often read, etc., is part of our CL knowledge about books and similar things, not of our lexical knowledge about "book".

A further general consideration argues against treating the interpretation of "finish NP" as a lexical process: such a treatment would leave a host of interpretations unaccounted for, and any plausible mechanism that would account for these would also account for the interpretations supposedly handled by the GL. Consider:

(24) Have you finished that piece of trash.

If the understander knows that the NP "that piece of trash" refers to a book, he will understand this sentence as he would if the NP was "that book". The interpretive possibilities are identical: i. e., in neutral contexts the activity of reading will come to mind as a plausible candidate, but any contextual information will override. What counts for the interpretation of "finish NP" is not the meaning of the NP, but its denotation: if the NP refers to a book, then a certain range of possibilities is available. That is, what counts is BOOK, not "book". Using the NP "the book" is a good way to activate BOOK, but there is an infinite set of NPs that will do the

job, such as "Charles Dicken's last effort", "The Generative Lexicon", "Pustejovsky, 1995", etc. A model of interpretation needs specifications (roughly) like those of the GL LR for "book" for all NPs that activate BOOK, so a natural design decision would be to eliminate the LR specifications and keep the CL specifications.

A striking characteristic of the LRs of Pustejovsky, 1995 is that many of their specifications are in fact predictable, and would not be allowed by any theory that eliminates redundant information. On the other hand, unpredictable specifications are often not included in the LRs. For instance, for "book" the specifications on p. 101 do not distinguish "book" from other lexemes it needs to be distinguished from, like "magazine", it is not clear how such information is to be specified in this framework.

Of more immediate relevance is that none of the information that is specified needs to be. In fact, not only should it not be specified in the LR of "book", but arguably it should not be specified for the concept BOOK: the fact that books are often, perhaps typically, read and written is common to books, magazines, reports, brochures, leaflets, newspapers, certain T-shirts, etc. It is a general property of things that bear writing that the writing is often meant to be read and was written. Thus a simple specification that a book is a physical object that typically bears textual information should allow such specifications to be inherited. The generative thing to do would be to generate the information from an LR (or CL representation) more like this:

book: physical_object_containing_text

There is a further problem with the GL approach: the LRs typically contain specifications that look a lot like words or complex expressions of English. For instance the LR of "book" contains *read, write, hold, information, physical object.* The problem is: what are these symbols? They can't be other lexemes, since each of these symbols is meant to be understood in a specific sense, and lexemes in the GL theory are merely a basis for the generative mechanisms that produce senses. Each of these words in fact has multiple senses, and **read** is no doubt to be understood not in the specific sense of reading aloud, or using eyes rather than fingertips, but in some more general way (not general enough to include reading diskettes or tea leaves). Do the symbols stand for senses, or concepts? If so, that means at least some individual senses must be stored somewhere in the system, in addition to being generated. To judge by the examples, there will be a large number of such distinct symbols, for all we know as many as there are generated senses. These symbols would need to be stored since they presumably participate in entailments that need to be specified. But if a good proportion of the senses used in specifying LRs have to be stored independently in some lexicon of primitives, it is not clear why we need a GL in the first place. At least, arguments appealing to

economy are weakened: lexemes could just point to individual senses, as in old-fashioned polysemy. How much does the GL save us having to store somewhere in the system ? We can't know until we know what the theoretical vocabulary of lexical description looks like, and where it comes from.

If we assume that the CL itself is generative, then the program of reducing polysemy can still be pursued by putting the entire generative burden on the CL. The problem just mentioned will then arise at that level, to which a traditional, often criticized answer is that concepts are made up of primitive conceptual elements, and entailments are stated in terms of primitives. I will not deal here with the interesting question of whether a GL-like approach, reformulated at the CL, offers help here.

In conclusion, one criticism of the GL approach is that it is overly lexical, and not generative enough, these two points being related: because it puts the generativity in the lexicon, it includes material in the lexical entry that doesn't have to be there, being derivable from more general CL knowledge, thus it is not sufficiently generative, and it is too lexical, because in many cases the lexicon is the wrong site for such specifications.

(Université Paris 8)

References

BERLIN (B.), KAY (P.)
1969, *Basic Color Terms: Their Universality and Evolution,* Berkeley, University of California Press.

BLOOM (P.), PETERSON (M.), NADEL (L.), GARRETT (M.)
1997, *Language and Space,* Cambridge, MIT Press.

CAREY (S.)
1985, *Conceptual Change in Childhood,* Cambridge, MIT Press.

CARTER (R.)
1980, "La Notion d'explication en sémantique", *Langue française,* n° 46.

FILLMORE (C.)
1982, "Frame Semantics", p. 111-138, *in Linguistics in the Morning Calm,* Linguistic Society of Korea, ed., Seoul, Hanshin.

FODOR (J.)
1994, *The Elm and the Expert,* Cambridge, MIT Press.

HAIMAN (J.)
1980, "Dictionaries and Encyclopedias", *Lingua,* n° 50, p. 329-357.

KATZ (J.), FODOR (J.)
1963, "The Structure of a Semantic Theory", *Language,* n° 39, p. 170-210.

LANGACKER (R.)
1987, *Foundations of Cognitive Grammar,* vol. I., Stanford University Press.

LEHRER (A.)
1974, *Semantic Fields and Lexical Structure,* Amsterdam, North-Holland.

LEHRER (A.), KITTAY (E.)
1992, *Frames, Fields, and Contrasts,* Hillsdale, Lawrence Erlbaum Associates.

JACKENDOFF (R.)
1985, "Multiple Subcategorizations and the q-Criterion: The Case of *Climb*", *Natural Language and Linguistic Theory,* n° 3, p. 271-295.

1990, *Semantic Structures,* Cambridge, MIT Press.

1997, *The Architecture of the Language Faculty,* Cambridge, MIT Press.

LAKOFF (G.)
1987, *Women, Fire and Dangerous Things : What Categories Reveal About the Mind,* University of Chicago Press.

LEECH (G.)
1969, *Towards a Semantic Description of English,* Bloomington, Indiana University Press.

PUSTEJOVSKY (J.)
1995, *The Generative Lexicon,* Cambridge, MIT Press.

TALMY (L.)
1972, *Semantic Structures in English and Atsugewi,* Doctoral Dissertation, University of California, Berkeley.

1985, "Lexicalization Patterns : Semantic Structure in Lexical Forms", *in Language Typology and Syntactic Description,* vol. 3, T. Shopen, ed., Cambridge University Press.

ZABLIT (P.)
1991, *Construction de l'interprétation temporelle en langue naturelle,* Thèse de doctorat, Université de Paris-Sud, Orsay.