Re-cycling in the Encyclopedia[1]

Richard Hudson

University College London, Department of Phonetics and Linguistics, Gower Street, London, WC1E 6BT, United Kingdom
E-mail: dick@linguistics.ucl.ac.uk

and

Jasper Holmes

Rijksuniversiteit Groningen (??), [affiliation + street address]
E-mail: uclyjho@ucl.ac.uk

0. Introduction

The principle of "recycling" holds that (mental) concepts are "recycled" rather than duplicated. On the one hand, a concept that is used outside language may be recycled as the meaning of a word (language-external recycling); and on the other, a concept which serves as the meaning of one word may be recycled as part of the meaning of another (language-internal recycling). The alternative in both cases is to postulate a concept which is both similar in "content" but distinct in status (and name). At one level, this is just a matter of common sense (Hudson 1985): why should we, as learners, build two distinct concepts when one concept could do both jobs? For ex-

[1] We would like to thank Joe Hilferty and an anonymous reviewer for comments on an earlier version of this paper.

ample, if a child has a concept "bicycle" for dealing with the world (of bicycles), why should it construct a distinct concept for dealing with the word *BICYCLE*? Much more likely, surely, is that a single concept will be used for both purposes. Similarly, if a child has a concept "bicycle" for dealing with the word *BICYCLE*, why should it construct a further concept to fill the vehicle role in the meaning of the verb *CYCLE*?

However, this principle is clearly not self-evident because there are well-established intellectual traditions which deny it. Theories of meaning traditionally distinguish sharply between dictionary-meaning and encyclopedic information, which at least allows the interpretation that the concepts involved must be different: there must be one "bicycle" concept defined by the dictionary, and a different one defined by the encyclopedia. Another tradition that ignores the principle is practical lexicography, where different words are defined separately (often by separate lexicographers) without any attempt to build one word's meaning into that of the other. E.g., in the Collins Cobuild English Language Dictionary we find the following two definitions:

(1a) **Skill** is the knowledge and ability that enables you to do something such as a job, game, or sport very well.

(1b) Someone who is **skilful** at doing something does it well.

Does this mean that skill requires a higher level of competence ("... very well") than merely being skilful ("... well"), or that it is possible to be skilful at doing something which is not covered by the list of examples mentioned for skill? If the principle of recycling had been respected, the second definition would have included the word *skill* (e.g. "Someone who is skilful has a skill"). In view of these departures from the principle, it clearly needs to be supported by evidence.

The theory of recycling allows testable predictions. If a concept is involved in the definitions of two words, it should carry the same uncertainties of interpretation in both words. This seems to be true in at least some cases. For example, the meaning of *MOTHER* applies clearly in some cases, but less clearly in others where the traditional roles are divided between a birth-mother and a nurture-mother (Lakoff 1987:74-84). As predicted by recycling, exactly the same uncertainties arise in interpreting words such as *PARENT*, *GRANDMOTHER* and all the terms whose definitions build on "parent" - *AUNT*, *UNCLE*, *COUSIN* and so on (Hudson 1995:63-72). This supports the view that the definition of (for example) "parent" really is 'mother or father', where "mother" is exactly the same concept that is also the sense of *MOTHER*.

The purpose of this paper is to explore the concept "cycling", the sense of the verb *CYCLE* as used in sentence (2) (which is normal British usage):[2]

[2] It appears to be less familiar in America, where the verb *BIKE* is more commonly used (cf. below).

(2) I cycled to work.

The discussion will provide further evidence for language-external recycling of concepts. If the same concepts which are linked to words are also linked to more general conceptual "frames", then in principle anything language users know about bicycles should be able to affect the linguistic behaviour of *CYCLE*.

However, our conclusion will not be that word meanings are simply non-linguistic concepts that have names. We shall show that the language itself imposes a structure of its own on the concepts, and argue that far from refuting the claim of recycling this is exactly what we should expect. The analysis will be couched in terms of Word Grammar (henceforth WG; cf. Hudson 1990, 1998).

1. Bicycles

The most obvious fact about cycling is that it typically involves a bicycle. (We return to atypical cases below.) Any analysis of Cycling must therefore recycle Bicycle.[3] It would be ridiculous to define the means of locomotion in cycling without any reference at all to Bicycle.

What, then, is a bicycle? This question has been answered at length by Wierzbicka (1985:104-123), whose discussion shows how much we all know about bicycles, and whose definition fills a whole page (ibid:112), of which we quote the first few lines:

(3) A KIND OF THING MADE BY PEOPLE,

it is made for one person to be able to go by means of it from one place to another

faster than by walking and with less effort...

The definition is expressed in part in terms of a universal metalanguage containing a limited number of terms. Wierzbicka claims that this "Natural Semantic Metalanguage" reflects a "language-like innate conceptual system", and could in principle be used as the sole metalanguage for semantic definitions (Wierzbicka 1996:22).

We find this claim unconvincing. First, we notice that her own definition of Bicycle applies the principle of recycling by using terms which are not part of the universal metalanguage. The passage quoted includes *faster*, *walking* and *effort*, and in the later parts of the definition we find *legs*, *feet*, *hands*, *wheels*, *frame*, and *stick out*.

[3] From now on we shall use upper case initials when referring to the names of concepts, to save repeating the phrase *the concept "X"*. Words will continue to be written using small capitals.

Without these terms, the definition would have been much worse - not only impossibly long and tedious, but also less revealing.

However, more importantly, there is a fundamental conflict between her theory and the principle of recycling. Her aim is to define all concepts in terms of the same set of primitives, but this excludes recycling in principle. As we have just seen, *WHEEL* is included in her definition as a concession to practicality, but this presumably counts as a weakness in the analysis, when judged by her principle, whereas we consider it a strength. When the principle is applied rigidly, the result is surely a poor analysis. To take a simple example, her definition of Bicycle repeatedly refers to the pedals, but avoids using the term *PEDAL* because this is not part of the universal metalanguage. Consider the following extracts from later in the definition:

(4a) it has two parts for the person's feet

which are attached to the frame near the bottom and which are connected with the wheels

so that by pushing these parts with one's feet one can cause the wheels to turn

(4b) the parts for the feet are as small as they can be without being too small to support a person's feet

The definition leaves everything to the reader's linguistic and pragmatic skills, whereas a good analysis would surely make the link between the "parts for the feet" explicit by assigning them a single name within this definition.

Worse still, the ban on recycling prevents the definition of one concept from building on that of another concept. What about Pedal itself? It too needs a definition, but this will have to start from scratch, rather than building on the information about pedals in the definition for Bicycle. This is exactly the opposite of what recycling demands. We believe that the definition of Bicycle should refer to Pedal, so that each definition supports the other; and the reason why we believe this is not just for analytical brevity, but because we believe that knowledge is, in fact, integrated in this way.

Another theoretical claim that underlies Wierzbicka's analysis of Bicycle is that a distinction can be drawn between the "mental dictionary" and the "mental encyclopedia" (1985:113, 1996:335). This allows her to exclude a great deal of historical information which she finds in the Encyclopedia Britannica, on the grounds that most people know what a bicycle is without knowing about its history. This may well be true (though most people probably know a few things about the history), but what does the claim mean? As we commented earlier, one possible interpretation is that there are in fact two concepts, one of which acts as the meaning of *BICYCLE* while the other is the locus of the encyclopedic knowledge. But if that is the case,

what is the relationship between these two concepts? Surely the whole point of the distinction is to distinguish different kinds of knowledge about the *same* concept? But if the same concept is involved in both kinds of knowledge (as we assume is the case), what does the distinction mean in psychological terms? After all, it is generally accepted that radically different kinds of knowledge may converge on a single concept - consider, for example, how different a lexeme's phonological structure is from its syntactic classification - so there is no general requirement that a concept's characteristics should be in some sense uniform. The issues that arise are too general and fundamental to pursue here, but we register serious doubts about this distinction.

So, what *is* a bicycle? As suggested for this very concept in Hudson (1995:24-28), we believe that Bicycle is part of a network of concepts that define each other, each concept being recycled in the definitions of others. As Wierzbicka's definition shows, the concepts relevant to Bicycle are of different types and have to do with at least the following: a bicycle's function, structure, mode of operation and size. We shall not attempt a "complete" definition simply because we do not believe this is a meaningful target. After all, if we are trying to model the knowledge of actual speakers, we must accept that different speakers have different amounts of knowledge.[4]

Here is an attempt at the beginnings of a prose definition:

(5a) A bicycle has a frame, two wheels, a saddle, pedals, a chain and handlebars.

(5b) A bicycle's rider sits on its saddle.

(5c) A bicycle's rider holds its handlebars.

(5d) A bicycle's rider pushes the pedals round.

(5e) A bicycle's rider rides it.

This definition contains much less information than Wierzbicka's, so it is not fair to compare its length with hers; but it is fair to point out how its length is minimised by the repetition of single terms like *PEDAL* which will be defined separately. In fact, all the statements in (5) form part of the definitions of the other terms involved: just as Riding is part of the definition of Bicycle, so Bicycle is part of the definition of Riding. Circularity is avoided because neither concept is defined exclusively in terms of the other. The idea that concepts are defined by their relationships to other concepts is reminiscent of traditional "field theories" of semantics, but unlike those theories

[4] E.g., Wierzbicka reports that spokes are mounted tangentially, i.e. at right angles to the radius of the hub. Some people know this, others do not. One of us learned it from Wierzbicka.

the notion of "contrast" plays no part here; a concept is defined by what it "is" (i.e. by its positive links to other concepts) rather than by its contrasts with what it is not.

To further clarify this idea, consider (5b). This statement helps to define Bicycle by specifying that a bicycle is ridden by someone who sits on its saddle. It also helps to define a number of other concepts: Rider (one kind of rider sits on a bicycle's saddle), Sitting (one kind of sitting is the relationship between a bicycle's rider and its saddle) and Saddle (one kind of saddle is the part of a bicycle where the rider sits). Ultimately, the analysis of any one concept is complete only when every other concept in the (same) mind has been analysed. We shall partially satisfy this obligation for Riding, but we shall have to leave all the other concepts undefined.

2. Networks as flexible frames

One of the leading ideas in recent theorising about lexical semantics has been the idea that concepts are linked in complex structures called "frames". These are the foundation of the theory of Frame Semantics (Fillmore 1985, Fillmore & Atkins 1992), according to which a word's meaning is understood "with reference to a structured background of experience, beliefs or practices, ... the background frames that motivate the concept that the word encodes" (Fillmore & Atkins 1992:77). The classic examples of frames are Commercial Transaction (background to Buying, Selling, Paying, Charging and so on) and Risk (background to the meanings of *DANGER*, *VENTURE*, *GAMBLE*, *RISKY* and other words, in addition to *RISK* itself). More generally, there are global theories of knowledge which claim that all knowledge is organised in terms of frames (Barsalou 1992).

One of the attractions of these ideas is the importance they give to relationships, in contrast with theories that merely recognise "associations" among concepts (what Barsalou calls "feature-list representations") - e.g. an association between Bird and Feathers, Wings, Eggs and so on. The structures we actually find are more like a bicycle frame, where each part has a definite (and rigid) relationship to each other part. Cognitive science has established that mere association is not enough, so we must pay attention to the specific relationships among concepts (Reisberg 1997:280-285). The essential characteristic of frames is that they define a set of more or less specific relationships among the concepts that they bring together - e.g. "buyer", "seller", "money" and "goods" in the frame for Commercial Transaction.

A more negative feature of frames is the implication that knowledge is divided into frame-sized packages, each with its own natural boundaries and identity. As Lehrer & Kittay (1992:16) point out, this raises serious analytical problems to which it is hard to see any solution: "How do we decide that terms belong to different semantic frames (...), rather than saying that they have different meanings within one frame

(...)?". For instance, by claiming that Money is part of the Commercial Transactions frame, are we denying that it is part of the Wealth frame, the Banking frame or the Work frame? Similarly, Goods seems to belong just as much to Ownership as to Commercial Transactions. When boundaries are problematic, it is often worth considering whether they really exist and to look for a theoretical alternative in which the boundaries concerned play no part.

In this case, the alternative is the widely held view that knowledge constitutes a *network* (Reisberg 1997:257-303). A network can give the same information about relationships as frames, but without demanding any boundaries between parcels of knowledge. The difference between networks and frames is largely a matter of metaphor. In a frame analysis each concept is a box, whereas in a network it is a point or node. In both cases the concept is defined by its relationships to other concepts, but in a frame analysis the definition is contained in the box, whereas in a network it is just the totality of links from the concept concerned to other concepts. In a frame analysis it is possible to imagine a distinction between the "defining" relationships, contained inside the box, and other, descriptive or encyclopedic relationships, which are left outside; but such a distinction is impossible, in principle, in a network. Frame analysis raises the problem of frame boundaries defined by Lehrer & Kittay; network analysis does not.

Returning now to the definition of Bicycle, we can show how the prose definition given in (5) can be translated into a network. Each of the concepts named is represented by a node, and each relationship is a function which is represented by a labelled arrow that points from the argument to its value. Some of the nodes are labelled for convenience, but the labels are actually redundant, since every node is already defined uniquely by its relations to other nodes. Suppose we assume, to start with, that each verb-noun pair in the prose definition corresponds to just one relationship in the diagram. The result is shown in fig. 1, which contains some relationships that are rather implausible, such as "sits on" and "pushes round".

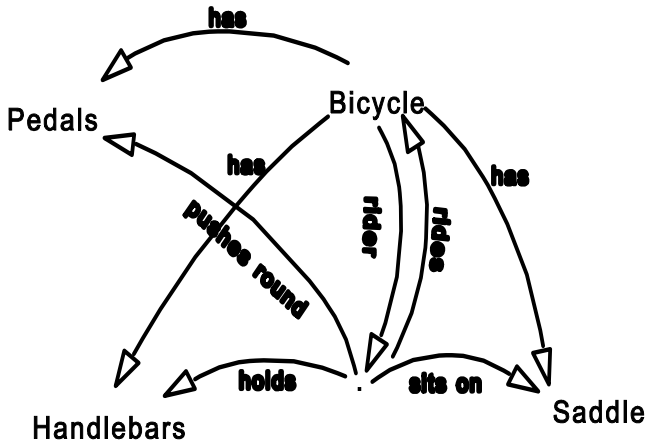

Figure 1

The diagram does not yet include all the information from (5): for example, it ignores the frame, the wheels and the chain. The main point, however, is that a concept may be shown as just one node on which several relationships converge, providing its "definition". Thus Bicycle is by definition the concept which stands in this particular set of relationships to Pedal, Handlebars, and so on, and the rider (shown by the unnamed dot) is defined by its relationships.

Various improvements to the network in fig. 1 are possible. One unsatisfactory feature is the naive way in which the distinction between concepts and relationships is applied, based as it is directly on grammatical categories: nouns define concepts and verbs define relationships. Like many others (e.g. Jackendoff 1983:67), we believe grammatical class to be irrelevant to status in conceptual structure, so concepts may be defined by verbs as well as by nouns. Accordingly, we assume that the network will include concepts for "states of affairs" like Sitting (or Sitting-on) and Pushing (or Pushing-round) which can be exploited here. This gives us the network in fig. 2.

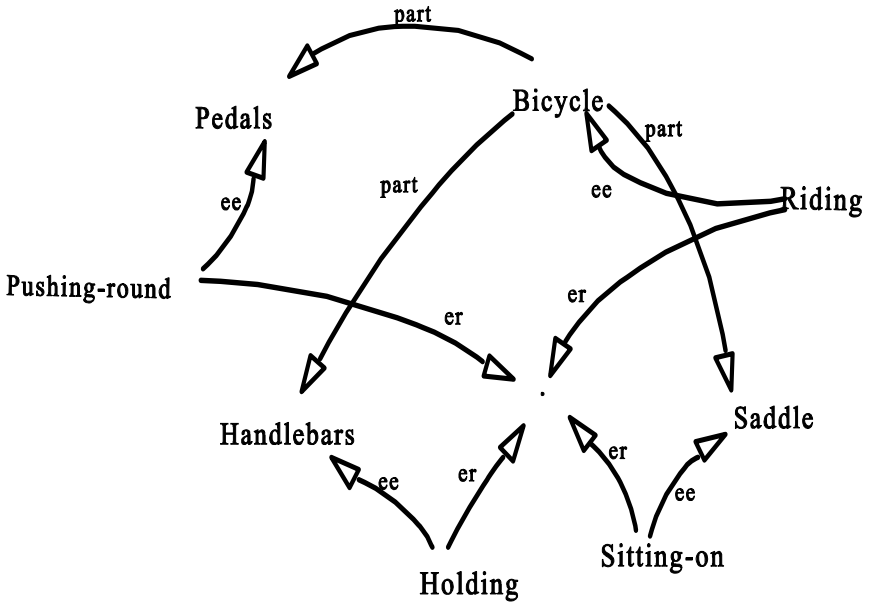

Figure 2

All the relationships are now rather general and drawn from a small list - in this case just Part (replacing the less informative Has), and the two provisional role names Er (for Sitter, Pusher and so on) and Ee. The diagram is more complicated than the first because single direct relationships (e.g. “sits on”) have been replaced by pairs of relationships (Er, Ee) to a linking concept (Sitting-on). However, this price is worth paying for the benefit of a more homogeneous semantics. Instead of assuming that nouns map onto concepts and verbs onto relationships, we can assume that both kinds map onto concepts; and instead of assuming that noun meanings are defined by links whereas verb meanings are links, we can define both in the same way. Thus Sitting is a concept node linked to others (Posture, Vertical, Bottom, Seat, Weight, On and so on), and just like Bicycle, the concept Sitting is defined by the sum total of these links.

The unified view of word meaning illustrated here is rather uncontroversial among cognitive linguists, who would probably agree for example that the concept Arriving can be expressed equally well either by a verb or by a noun (*ARRIVE*, *ARRIVAL*) and most relationships double up as the sense of some noun (*DIFFERENCE*, *MIDDLE*, *RELATIONSHIP* and so on). Somewhat more controversially, however, we are not convinced by those like Wierzbicka (1988) and Langacker (1987) who argue that grammatical differences always indicate subtle semantic differences of “construal” which should affect the semantic structure. In our view, the sense of the verb *ARRIVE* is precisely the same concept as the sense of the noun *ARRIVAL*; they are exact synonyms

(though semantic differences result, of course, from the effects of verb and noun inflections).

Our approach is therefore basically orthodox in relation to the tradition of cognitive linguistics, but it departs from the tradition of predicate logic by denying the simple mapping between syntax and semantics which underlies much of predicate logic; in this, however, it follows a widely accepted recent tradition which includes ACT-R (Anderson 1993), Conceptual Dependency analysis (Schank & Rieger 1974) and Conceptual Graphs (Sowa 1984). The assumption in predicate logic is that the semantic predicate is defined by the syntactic predicate (which includes at least the verb), and its arguments are defined by the syntactic subject and possibly other syntactic "arguments". For instance, *Pat sat on the bicycle* maps to something like "Sitting-on (Pat, The-bicycle)". The predicate and its arguments have quite different statuses in the logic - for example, the arguments may be variables, but this is (normally) not allowed for the predicate. In our analysis, in contrast, the semantic structure consists of three concepts (Pat, The-bicycle, Sitting-on) and two relationships (Er and Ee), each of which has an argument (Sitting-on) and a value (respectively Pat and The-bicycle). The two analyses are contrasted in fig. 3. In spite of this fundamental difference, our analysis does preserve one of the assumptions of predicate logic, which is that concepts are not linked as equals (except in coordination). Wherever a pair of concepts is linked, the link "belongs" to one of them; for example, the Er link belongs to Sitting-on, not to the cyclist, because we know that Sitting-on needs an Er, a sitter (just as a grin needs a grinner, in spite of the Cheshire Cat); but a person need not be involved in Sitting-on, and only becomes a sitter by virtue of a link to Sitting-on. In the diagram this asymmetry is shown by the direction of the arrow, which goes from the owning concept (its argument) to the other concept, which is its value.

PREDICATE CALCULUS

Sitting-on (Pat, the-bicycle)

WORD GRAMMAR

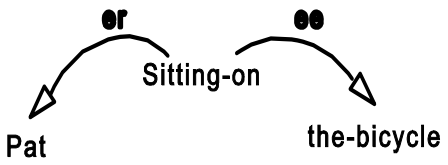

Figure 3

In summary, then, grammatical word classes are not reflected in the semantics. As far as verbs and nouns are concerned, they are all linked permanently to at least one sense, and all senses have the same status: they are concepts which are linked in many different ways to other concepts. Instead of being divided into "frames", knowledge spreads without boundaries across the whole network. We must now address an important question for any network theory, which is how to distinguish the various links from one another.

The example just given might suggest that we are committed to searching for a very small set of primitive relationships such as Er and Ee, but this is not so. Indeed, as we shall show below, it is possible to interpret the label *er* either as the name of a constant (like the traditional Agent) or as a variable whose value depends on the owning concept - the rider of Riding, the sitter of Sitting and so on. As in Sowa's theory of Conceptual Graphs (conveniently summarised in Luger & Stubblefield 1993:368-378), we assume that a large number of relationships may be distinguished: rider, sitter, colour, name, meaning, pronunciation, and so on. This is helpful, even essential, if each concept is linked in many different ways to other concepts. For example, a word has a pronunciation, a spelling, a meaning, a word-class, a language, a style and perhaps other attributes, each of which is shown by a separate link, and each of these links must be distinguished from the others.

It is controversial to assume an open-ended list of relationships, because each one must somehow be defined; the same objection can be made of course to any set of concepts, but it is easier to see how non-relational concepts define each other. It is often objected that the attractions of a network approach presuppose a small, predefined list of relationships (Reisberg 1997:280-281), so we must ask how relation-

ships are defined. The answer is that they are defined in just the same way as concepts: by their place in the total network, so the labels are strictly redundant. Firstly, a complex relationship can be defined in terms of simpler ones: for example, the Grandparent relationship can be defined in terms of Parent (Grandparent of X = Parent of Parent of X), and Rider-of (a direct link between the rider and the vehicle) in terms of Er, Ee and Riding, as shown in fig. 4.

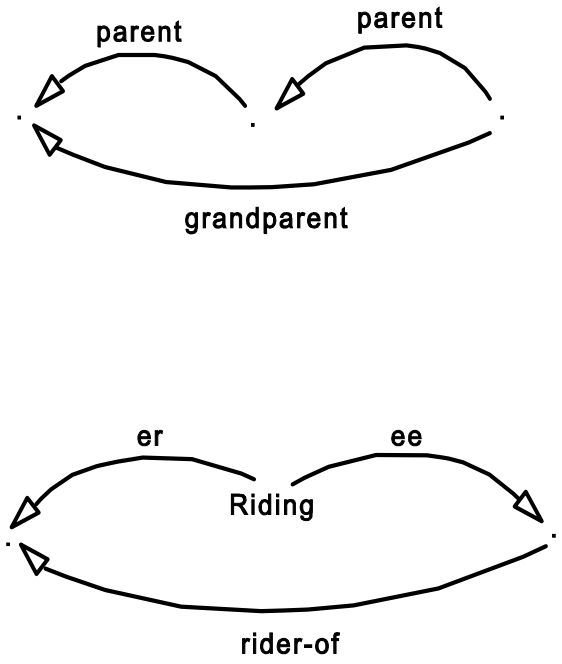

Figure 4

The second way of defining relationships in the network is the same mechanism by which concepts are linked in an "Isa" hierarchy. So just as Dog isa Mammal, we can classify relationships: Parent isa Relative, Rider isa Er, and Rider-of isa User-of. In WG diagrams the Isa link is shown by a small triangle which rests on the supercategory and whose apex is linked by a line to the instance. Fig. 5 shows how this system can also be applied to relationships. As far as the analysis of Bicycle is concerned, this means that we can define Rider as a derived relationship, based on the concept Riding and the two more basic relationships Er and Ee, and similarly for Pedaller, Holder and Pusher. If necessary we can even define specific variations on the Part relationship, such as "Saddle-of" and "Frame-of", in terms of the basic Part relationship plus the categories to which the things concerned belong.

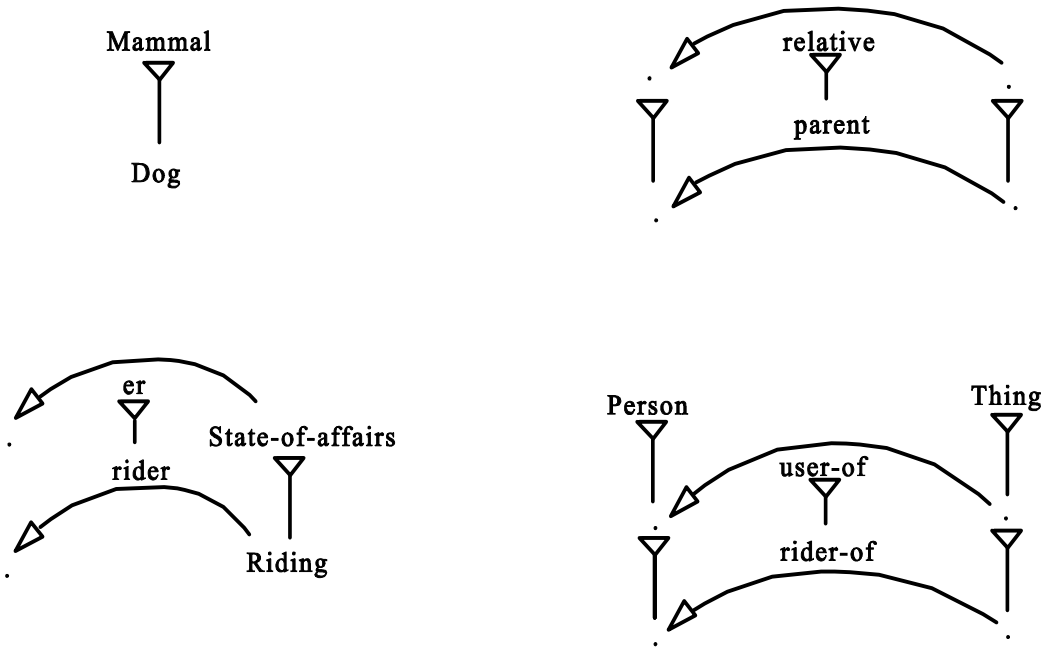

Figure 5

Such derived relationships may be needed in natural-language semantics for possessive constructions (e.g., *my nose* involves a relationship Nose-of), but they also have two advantages in network analysis. First, they reduce long-distance relationships to combinations of local relationships; for example, Grandparent establishes a direct link between nodes which would otherwise be linked only indirectly, but once established this link may itself be used in defining other relationships (e.g. in this case, Great-grandchild of X = Child of Grandchild of X). This has a bearing on the shape of the network and seriously affects the way in which activation will spread from node to node when the network is in use. Clearly the process of defining one relationship in terms of a chain of others is selective, because spreading activation is selective; so it is important to discover the principles on which the selection is based.

The second advantage of recognising derived relationships is that they combine the flexibility of a network with the internal specificity of a frame. Suppose a concept has many different parts whose internal relationships are rigidly fixed; this internal structure can easily be defined in a frame, but it can also be specified in a network provided that the various parts can be distinguished from one another in terms of their relationship to the whole - in other words, if they are defined "functionally". This is made possible by derived relationships. Take the example Bicycle. If a network for Bicycle defines the functions Frame-of, Front-Wheel-of and so on, it can then define the relationships among their values, as in fig. 6, which shows the spatial relationships among some of the parts. In this diagram, we use the two dimensions

as an analog equivalent of true space, just as one might use the left-right dimension to represent time; but unlike the left-right convention we assume that this picture is in fact a plausible approximation to our actual mental representation of Bicycle (some kind of mental image, rather than a truly propositional structure - "a pattern in the 2½-D sketch that is loaded from long-term memory rather than from the eyes"; Pinker 1997:286). It should be obvious how this network could be expanded to include the remaining parts.

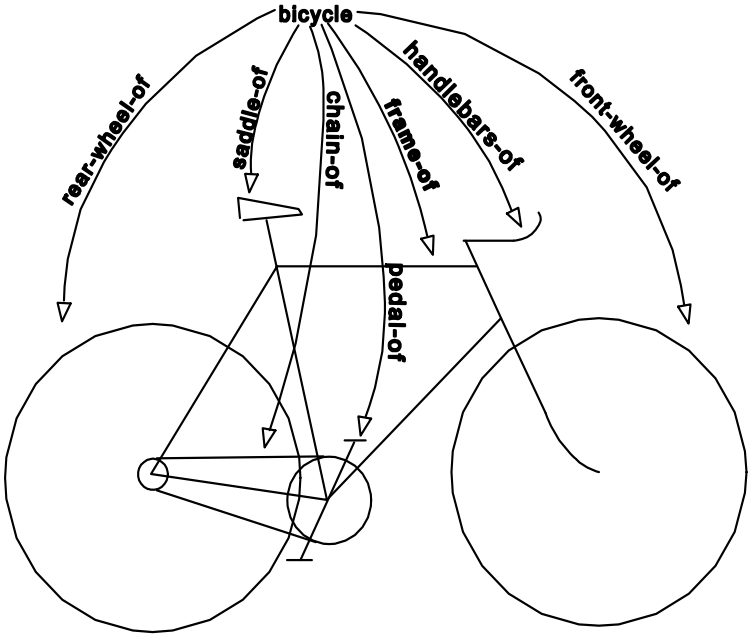

Figure 6

The analysis could be taken much further; for example, we could build an "Isa" hierarchy for all the objects (Bicycle, Pedal, etc), people (Rider) and activities (e.g. Riding, Holding), and provide fuller definitions of the parts in terms of their physical, interactional (e.g. force-dynamic) and functional properties. We could also discuss the status of non-typical bicycles by discussing the logic of default inheritance and the Best Fit Principle, which are central to WG theory. However, the main point that we have established is that the concept Bicycle can be fully defined by its relationships to other concepts, and that these relationships integrate the concept into a vast network which includes the totality of our conceptual knowledge. We shall now take this view for granted and move on to the definition of cycling.

3. Cycling as riding

The meaning of *CYCLE* is closely related to that of *RIDE*, so *ride a bicycle* is a close paraphrase. However, the two are not exact synonyms. The differences shall be explored in a later section; meanwhile, our task is to reveal the similarities.

If the senses of *CYCLE* and *RIDE* are the concepts Cycling and Riding, the first step is to show that Cycling is a kind of Riding - that Cycling isa Riding, as shown in fig. 7, which also includes an Ee link (for the "ride-ee", the thing ridden) whose value isa Bicycle. In short, the diagram shows that cycling is riding a bicycle.

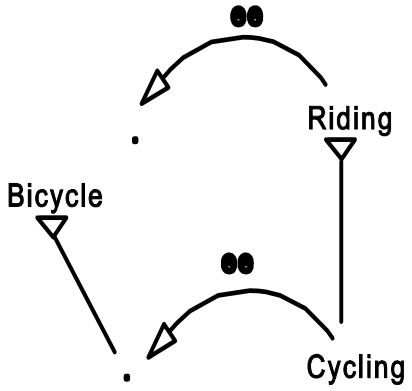

Figure 7

Of course, we know far more about Riding than the fact that it has an Ee. Riding is the sense of *RIDE* in examples like (6a) below:

(6a) I rode the bicycle/*car/*plane/*boat.

(6b) I drove the car/*bicycle/*plane/*boat.

(6c) I flew the plane/*bicycle/*car/*boat.

(6d) I sailed the boat/*plane/*bicycle/*car.

Clearly, Riding is restricted to certain kinds of vehicles - bicycles, but not cars or planes. Riding is not limited to bicycles, however: it is also possible with motorcycles and horses; indeed, horses are probably the default, as can be seen from the normal meaning of *RIDE* when used intransitively (e.g. *I go riding every evening*). What these three "vehicles" have in common is that the rider sits *on* them, not *in* them, so this is another fact that we can include: the rider is in the "on" relationship to the vehicle.

Thirdly, the rider is in control of the vehicle, in contrast with a different, intransitive, use of *RIDE* (followed by *in* or *on*), which we can call Mere-Riding. Though the preposition varies with the vehicle, *any* vehicle is possible in (7b) and (7c). This is further confirmation of the fact that a different sense is involved.

(7a) I rode the bicycle/*bus/*ship/*train.

(7b) I rode on the bicycle/bus/ship/train.

(7c) I rode in the car/plane/boat.

If A rides on the back of B's bicycle, A is riding *on* the bicycle, but not riding it. If Riding involves being in control, it must be an example of Controlling, so Riding isa Controlling (as Cycling isa Riding). It must also isa Mere-Riding, because whenever you ride a bicycle you necessarily ride *on* it as well (and likewise for horses and motorcycles). Consequently, Riding isa Mere-Riding as well as Controlling: if you ride a bicycle, you ride *on* it and control it. All these relationships are shown in fig. 8.

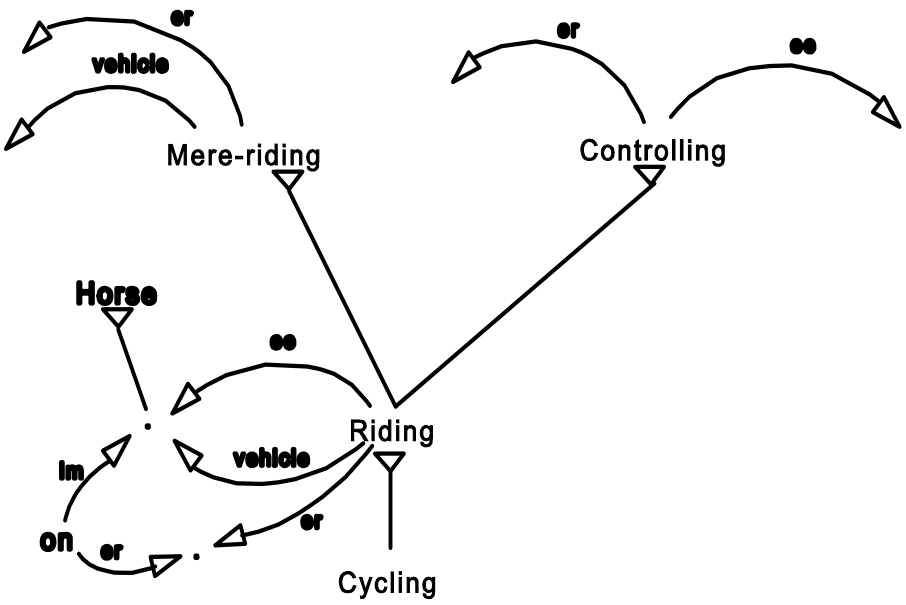

Figure 8 (the relationship labelled *lm* is Langacker's 1987 "landmark")

4. Language and concept formation

With part of the analysis in place, we can pause to consider a very general question: what role does language play in the learning of concepts? The principle of recycling predicts that language and general thinking share the same concepts: Bicycle is used not only in understanding the word *BICYCLE* but also in coping with bicycles in daily life. The question is whether these concepts are learned through language or through direct experience. Our most honest answer is that we (the authors) simply do not know, and cannot offer a well-developed theory of concept formation. However, some of the facts about the verb *CYCLE* are suggestive, and support the conclusion that some word meanings must be learned primarily on the basis of language, whereas others are probably learned primarily through direct observation. This conclusion is uncontroversial, but the relevant facts strike us as interesting.

It is easy to learn concepts for animals, foods and daily activities such as waking up, eating and walking simply by direct observation because the categories are clearly distinguished by bundles of cooccurring observable features (Rosch 1978:28). If concepts are defined by their links to other concepts, then a new concept is defined as soon as a unique bundle of links is established. It is likely that children learn the concept Bicycle in this way, since bicycles are so distinctive both visually and functionally. Similarly for Cycling. There is no reason to believe that either depends on the children's experience of the relevant word, provided they have direct experience of bicycles and of people cycling. The same is not true, however, of the concepts Riding, Driving, Sailing and Flying as illustrated in (6). Take Riding, which is applied to horses, bicycles and motorcycles:

(8) I rode my horse/bicycle/motorcycle/*car/*boat/*plane.

This list of vehicles is motivated by the fact that in each case the rider sits *on* the vehicle, not *in* it, and the vehicle goes on land, not water or air. It seems, then, that the verb *RIDE*, at least as used in examples like (8), always has the same concept as its sense. Similar conclusions apply to the other verbs, so our minds must also contain the concepts Sailing and Flying. But how did we learn these concepts? It is possible that we might have arrived at this classification of "transported going" on the basis of direct observation, but many other classifications are equally easy to imagine: e.g., we might have made a fundamental distinction according to whether or not the vehicle had wheels, or a motor, each of which would have been different from the one that we actually have.[5] Nor can we assume that each concept is defined by the skill

[5] The criterion of wheels puts bicycles with cars, not with horses, and that of a motor puts motorcycles with cars, not with bicycles.

that it demands - riding a motorcycle requires some of the skills of driving and some of riding, and an ability to ride a horse scarcely generalises at all to riding a bicycle.

In short, the four-way division into Riding, Driving, Sailing and Flying is not "natural" and determined by the way the world is. Rather, we assume that it is learned largely on the basis of language: when we learned the verb *RIDE*, we created a concept for its sense on the basis of how we heard this verb used. This is a coherent concept definable by a single bundle of links to other concepts, but it is not an inevitable outcome of non-verbal experience. It is learned on the basis of experience, but verbal experience plays a crucial part, so our concept formation has been guided by the idiosyncracies of English. We assume that these in turn reflect the history of English society, going back to the days when Riding was exclusively tied to horseback, Driving to horse-drawn vehicles and Sailing to boats with sails and when Flying was defined by birds and insects. If this conclusion is correct, we might expect other languages to have organised this part of experience differently - as is indeed the case. E.g., in German, the verb used with bicycles is the same as for cars (*FAHREN*), and different from the one used with horses (*REITEN*).

Examples like these seem to us to support strongly the view that "knowledge of the world" and "knowledge of language" merge in a single network of knowledge to which both linguistic and non-linguistic experience can contribute. It is true that some concepts are relatively independent of language while others are relatively dependent on it, but there is no clear boundary between them, and no clear advantage in trying to draw one. We shall see further evidence for this view in the following sections.

5. Cycling as a manner of motion

Mere-riding in general (and cycling in particular) is a way of moving, so the analysis must relate it to the notion of moving. In this section, we look more closely at the concept Moving, and the three relationships that define it. All appear in fig. 9.

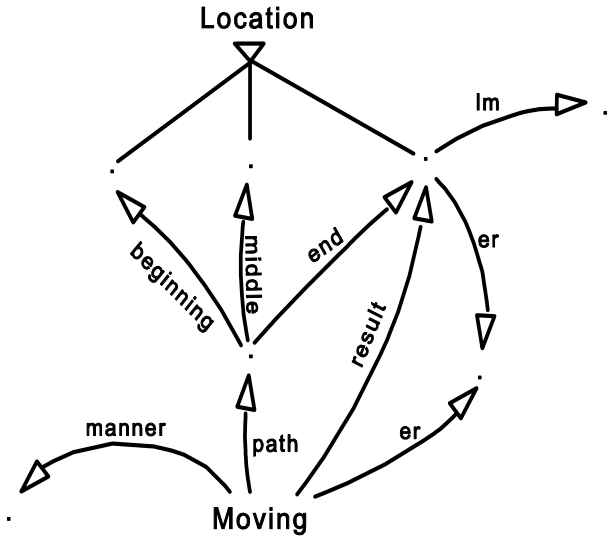

Figure 9

The first two relationships, viz. manner and result, are inherited from higher categories: any event can have a manner, and many events other than moving may have results. Specific to Moving is the fact that there is a path, the end of which corresponds to the result. The path is a spatial relationship between three parts: a beginning, a middle and an end. Each of these isa Location: a relationship between a located (er) and a place (lm). The last part of the path, its end, is the same as the result of Moving, which is also a Location. Though the path and the result are closely connected as described, they must nevertheless remain separate. In (9), the second adjunct refers to the duration of the result; in (10), it refers to the movement's duration along the path.

(9) We went to Manchester for two days.

(10) We went over the moors for two days.

Different kinds of Moving may be specified in terms of a *single* relationship, or in terms of a *combination*. (9) and (10) show that Going has a path as well as a result. Arriving is defined only in terms of its result (At), Rising in terms of its path (the end is Higher-than the beginning), Walking in terms of its manner (using the feet), Soaring in terms of both manner (rapid, uncontrolled) and path (as for Rising), Scramming in terms of both manner (rapid) and result (not here). A particular specification may be further detailed or overridden by the referents of syntactic dependents:

(11) We arrived in Manchester.

(12) He fell up the stairs.

(13) My uncle can walk on his hands.

In (11) the result is further detailed, in (12) the default direction is reversed, and in (13) the default manner is overridden. This is possible only if the overriding specification is close to the default value. Walking on your hands counts as walking under the Best-Fit principle (Walking is the closest kind of locomotion available in the lexicon). Moving around on your belly counts as a kind of Crawling for the same reason.

Where a kind of Moving is *not* specified as to one of the relationships, that relationship is still inherited and can be expressed in syntactic structure. In (14) the result verb *ARRIVE* has path and manner adjuncts; in (15) the path verb *RISE* has manner and result adjuncts; and in (16) the manner verb *CYCLE* has path and result adjuncts.

(14) We arrived by bicycle via Saddleworth.

(15) The float suddenly rose out of the water.

(16) We cycled over the moors to Manchester.

However, only those verbs that have a (compatible) lexically specified result may appear with certain result dependents:

(17) We went in the park.

(18) The float rose in the water.

(19) We cycled in the sand pit.

Going has a result, identified by the referent of the preposition in (17). Rising and Cycling, by contrast, have no specific result, so the prepositions in (18) and (19) refer to the location of the respective activities. In order to specify the result of Rising, we must exploit the path relationship, by using a preposition like *INTO*. *INTO*, since it refers unambiguously to a path, can also force a directional interpretation onto *CYCLE*.

The above differences lie behind the widely recognised contrast between manner-of-motion and direction-of-motion verbs (Levin 1993:263-264, Slobin 1996). *CYCLE* is a manner-of-motion verb, Mere-riding a concept that defines motion in terms of manner rather than direction. The manner shared by all examples of Mere-riding is that the rider is carried by a vehicle; in contrast with (say) a load of coal, what the vehicle carries is a person. Thus, our definition of Mere-riding must relate it to the

concepts Carrying, Vehicle and Person via the relationship Manner. Fig. 10 shows the details.

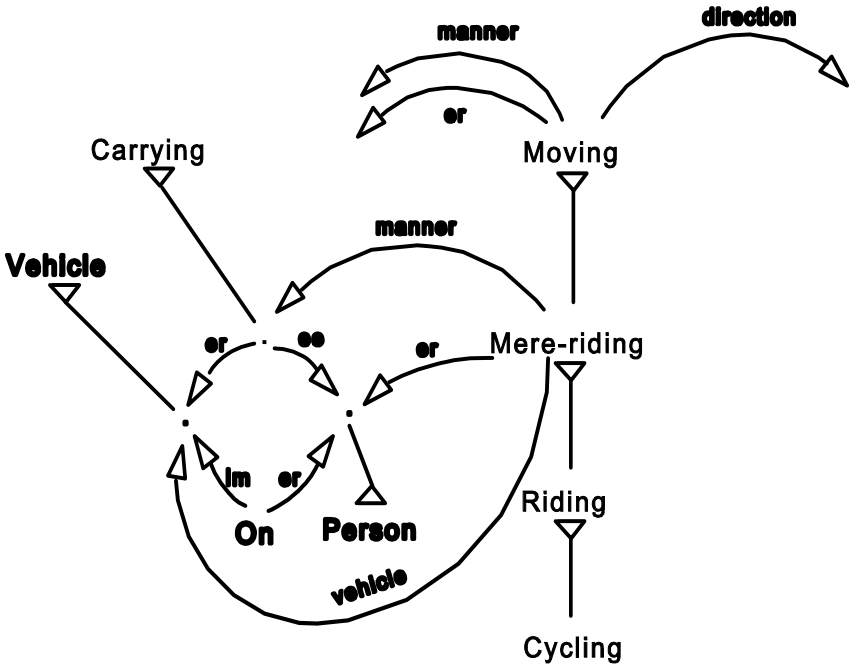

Figure 10

The network in fig. 10 could be paraphrased as follows: “You ride by being carried by a vehicle” or “When you ride you are carried by a vehicle”. Notice how the role relations are reversed between carrying and riding, since the ride-er is the carry-ee, and the carry-er is the location of the ride-er. This example justifies treating the movement and its manner as two different events, in spite of the fact that they are simply different conceptualisations of the same event.

We could now go on to define Carrying as a combination of supporting and movement; for example, if you carry a tray, you support the tray and move at the same time, so the carrying isa both supporting and moving. However the analyses given so far already show how the concept can be defined “upwards” in relation to more general concepts. These higher-level concepts are relevant because of default inheritance, so we now have a rich definition of Cycling which either by stipulation or by inheritance has the following characteristics:

(20a) Cycling isa Riding.

(20b) Cycling isa Mere-riding.

(20c) Cycling isa Moving.

(20d) The manner of Cycling isa Carrying.

(20e) The cycle-er of Cycling isa Person.

(20f) The vehicle of Cycling isa Bicycle.

(20g) The vehicle of Cycling is the carry-er of its manner.

(20h) The cycle-er of Cycling is the carry-ee of its manner.

(20i) The cycle-er of Cycling is the Er of the On state whose landmark is its vehicle.

These facts are all represented formally in the network, so the prose versions are a direct translation of the network notation. The semi-formal metalanguage reads oddly, but should be comprehensible and can easily be turned into ordinary prose (with the technical terms and concept-names retained):

(21a) Cycling is a kind of Riding.

(21b) Cycling is a kind of Mere-riding.

(21c) Cycling is a kind of Moving.

(21d) Cycling involves carrying.

(21e) A cyclist is a person.

(21f) The vehicle used in cycling is a bicycle.

(21g) In the carrying involved in cycling, the bicycle does the carrying.

(21h) In this carrying, it is the cyclist that is carried.

(21i) The cyclist rides on the bicycle.

We shall now consider cycling from a different perspective.

6. Cycling as pedalling

One of the special features of cycling, as an example of riding, is that it is the only kind of riding in which the rider has to provide the energy.[6] A closely related feature is that it is only bicycles that have pedals; and the link between the features is that the pedals receive the rider's energy and help to turn it into forward movement.

Somewhat surprisingly perhaps, these characteristics of cycling may be reflected in the lexical relations of the verb *CYCLE*. The verb has the same stem as the noun, whose sense is Bicycle: thus (for those who have *CYCLE* meaning Bicycle) to cycle is to ride a cycle. The same is true of the lexemes *BICYCLE* and *BIKE*: to bicycle is to ride a bicycle, and to bike is to ride a bike. The same is not true, however, of other verbs of riding. For horses, we have no verb *HORSE*, nor, for Mere-Riding, do we have verbs derived from *CAR*, *WAGON*, *TRAIN*, *PLANE* or any other kind of vehicle. Levin (1993:267-268) claims that most vehicle-nouns can be used as verbs, quoting examples such as *BALLOON* and *BOAT*, but our intuitions differ from hers, and we feel sure that even she could not use *CAR* or *AUTOMOBILE* as a verb. Admittedly, we can go boating (a pattern we shall consider below), but we can hardly boat to the next village. Even *MOTORCYCLE* and *MOTORBIKE* have no verb: you can cycle to the next village, but not motorcycle there. In short it is only verbs whose sense is Cycling that are "zero-derived" from a noun which names the vehicle.

On the other hand, there is a different area of vocabulary where this zero-derivation is very productive, and which may have provided the model for these exceptional lexemes. This is the semantic field which includes skating and skiing: when you skate you wear skates and when you ski you wear skis. The productivity of the system can be seen in neologisms like *ROLLER-BLADE*, which can be used either as a verb or as a noun. What these activities have in common is that the user provides the energy for forward movement through some kind of device fixed to their feet - let us call it (for want of an established term) "Foot-Powering". Admittedly, not all noun-verb pairs that define ways of moving fit this formula: exceptions include *SKI-BOARD* and *SURF-BOARD*, where the energy comes from elsewhere, and *SLEDGE* and *TOBOGGAN*, where there is not even a device fixed to the feet. Nevertheless, the pattern may be clear enough to motivate the zero-derivation of Cycling verbs.

However suggestive it may be, though, the linguistic evidence does not show that Cycling isa Foot-Powering, even if there are enough similarities to justify the zero-derivation. The pattern is productive in skating and skiing, but not in cycling. We can ride a tandem, but when doing so we are not tandeming, and likewise for *TRICYCLE* and *MONOCYCLE*. The most we can assume is that Cycling shares some fea-

[6] Horse-riding may be energetic, but the energy produced by the rider is not what keeps the horse and rider moving forwards.

tures with Foot-Powering which motivated the extension of the zero-derivation pattern - in the same way that it may perhaps have extended from skiing to ski-boarding and thence to surf-boarding, and from skiing to sledging.

Much more persuasive, however, is the analogy of a different verb in the field of cycling: *PEDAL*. This makes a close fit for Foot-Powering, both linguistically (to pedal is to use pedals) and conceptually, since it is the transmission of forward energy via a device fixed to the feet.[7] We can therefore conclude that the "pedalling" part of cycling isa Foot-Powering.

On a theoretical note, this is a clear example of the benefits of basing the analysis on a network rather than on frames. The analysis just suggested would be hard to express in frames because the frames for Riding and Foot-Powering are in general distinct, but show this small area of overlap. The frames cannot be simply combined by unification, but neither can either of them be treated as part of the other.

How, then, can we integrate Pedalling and Foot-Powering into the network for Cycling? The force-dynamic chain takes the energy from the rider's foot through the pedal (and chain) to the wheels, which convert the energy into forward movement. It is the forward movement of the bicycle that takes the rider forwards, so pedalling is only indirectly related to the rider's forward movement. Pedalling is not the manner of the cycling itself, but of its manner - it is the manner of movement for the bicycle, not for the cyclist. This is a happy conclusion, because Pedalling and Carrying would otherwise have been in competition as the value for the Manner function of a single concept. The solution, therefore, is to analyse Cycling as a kind of Riding whose manner is an example of Carrying whose manner is Pedalling.

The first step in the analysis is to focus on Pedalling. Unfortunately, as we have seen, this is itself a manner-of-motion like Skating and Skiing; so it has a manner of its own, which we shall call simply Pushing a device with the foot. This is the meaning of the verb *PEDAL*, as in (22).

(22) We pedalled across the lake in the pedal-boat.

The analysis of Pedalling is shown in fig. 11 as one example of Foot-Powering contrasting with (for example) Skating.

[7] The only uncertainty is whether pedals count as "fixed" to one's feet.

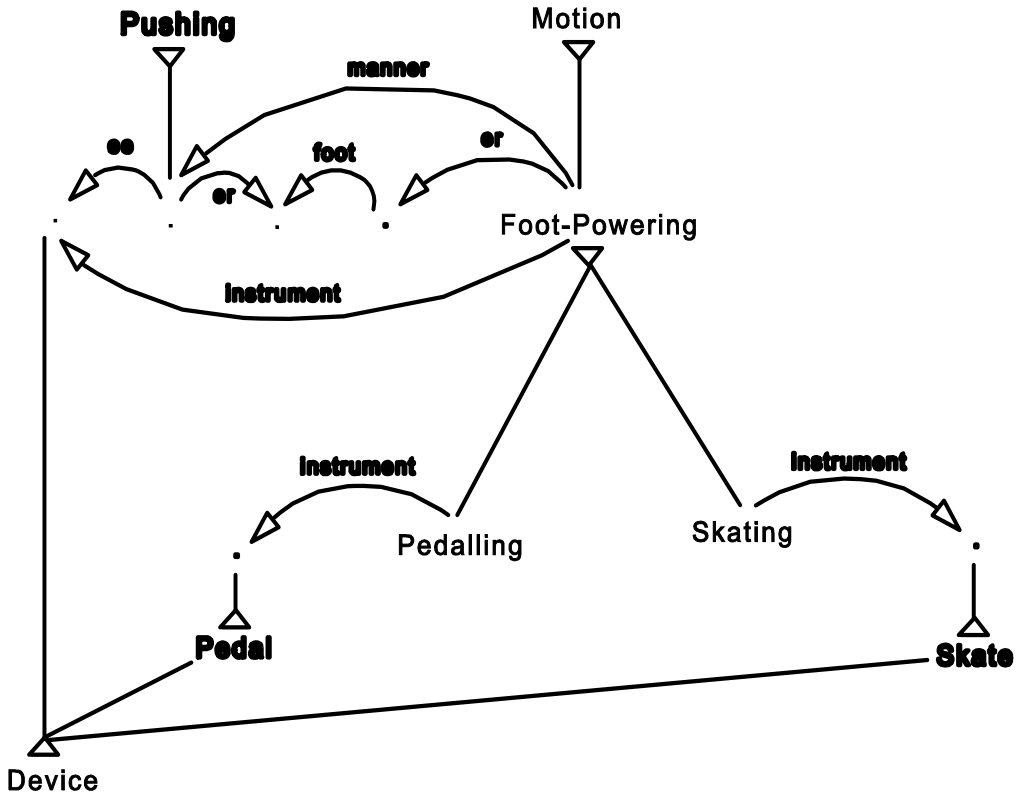

Figure 11

The next (and last) step is to show how the analysis of Pedalling is recycled as part of the analysis of Cycling. Ignoring the details that have already been covered in other diagrams, fig. 12 shows the overall structure of Cycling. It shows that pedalling is an essential component of cycling. This seems to us correct; for example, if we had to describe someone free-wheeling down a hill we would use *RIDE* (or a verb such as *COAST*) rather than *CYCLE*.

(23) I rode/*cycled down the steep hill without pedalling.

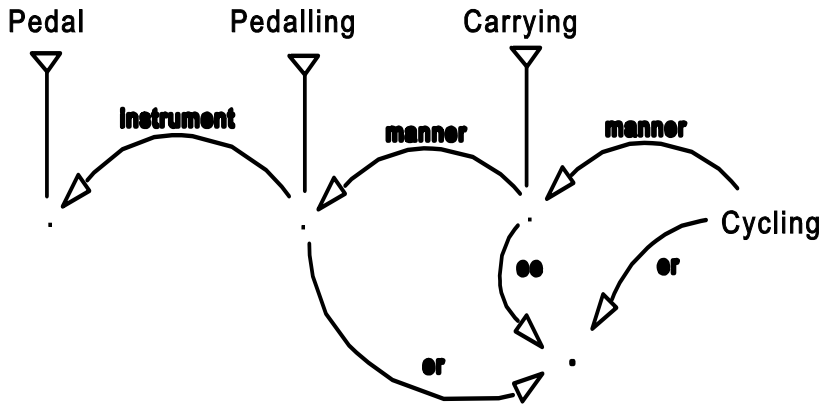

Figure 12

Pedalling is so important for Cycling that we may even be able to apply *CYCLE* to stationary exercise-bikes, where movement is obviously not relevant.

(24) ?I cycle on the exercise-bike for half an hour every morning.

However, if this is possible, it must be a special use of *CYCLE* where the default Moving is overridden - depending on experience and entrenchment, either a nonce exception or a lexicalised sub-case of Cycle. The “cyclist” is still supported by the bike, so half of the definition of Carrying is still relevant, but it is not complete Carrying, nor is it complete Cycling.

What this section has achieved is to show that Pedalling is an important element of Cycling, which can be incorporated in a network definition in spite of the complexities caused by the fact that Pedalling defines the manner of the manner of Cycling.

7. Cycling as outdoor fun and indoor exercise

The analysis so far has focussed only on the observable, physical aspects of Cycling - what the cyclist does, and what the results are in terms of movement. However we have ignored an important part of the activity, its purpose. Why do people cycle? The fact is, of course, that an infinite number of purposes are possible, and an ordinary sentence containing *CYCLE* leaves the choice completely open, so any of the continuations in (25) are possible:

(24) I cycled to work

a. to save time.

b. to save money.

c. to keep fit.

d. to remind myself of my boyhood.

e. to outwit the police.

f. to win a bet.

g. to avoid the Martians.

However, there are two purposes which have been picked out for special linguistic treatment: fun and fitness. At least two grammatical patterns signal that an activity is being done for such purposes. We therefore assume that they belong together conceptually as different manifestations of a single super-purpose. In the absence of an established term we shall call it "Life-Enhancing", the point being that the activity is done for the benefit it does to the actor's emotional or physical well-being, and not for the specific results produced by the action.[8] In the case of cycling, this means that the cycling is done for fun or to improve fitness, rather than in order to get to the destination - indeed, there may not be a destination (as when cycling round a race-track). The effect of these grammatical patterns, then, is to signal explicitly that what is normally just a by-product of the activity has been promoted to its main purpose.

The purpose of Life-Enhancing can be signalled by means of two different kinds of nominalisation. Both patterns are highly productive, so we shall start by exploring their use in domains other than cycling. One involves (once again) zero-derivation, but in this case the derivation of a noun from a verb. For example, *WALK* may be used as an ordinary common noun:

(26a) I had a walk.

(26b) I went for a walk.

(26c) The walk did me good.

(26d) I told them about my walk in the park.

(26e) What you need is a good walk.

If you walk to the kitchen to make a cup of coffee, you have not had a walk, less still gone for a walk. A walk is certainly an example of Walking, but it is walking whose purpose is not to reach a specific destination, but enjoyment. If you are invited to go for a walk, the route is open for negotiation and the destination may well be the same as the starting point. These observations suggest that the noun *WALK* has its purpose specified as Life-Enhancing. We do not know whether the verb has a default purpose

[8] It would be interesting to speculate about the role of language in the development of this concept.

(reaching the destination) or a completely unspecified purpose, though the result may be taken as the purpose in the absence of any other specification (it is the default).

The same seems to be true of all manner-of-motion verbs. A drive in the countryside aims at pleasure, not achievement, and likewise for a swim, a ride, a paddle, etc. However, complications arise with some verbs, one of which happens to be *CYCLE* and its synonyms. We cannot go for a bike or a bicycle; and presumably the reason for this is that as we pointed out earlier, the nouns *BIKE* and *BICYCLE* are already in use as the concrete nouns from which the verbs are derived - a bike/bicycle is what we ride on when we bike/bicycle.[9] Instead, we use a completely different noun, *BIKE-RIDE*. A bike-ride is certainly something one does for pleasure or fitness, and not just as a way to reach the destination: so when you cycle to work, you have not had a bike-ride. The choice of noun is not arbitrary, of course, if Cycling isa Riding, and *BIKE-RIDE* sits comfortably alongside *HORSE-RIDE*, *DONKEY-RIDE* and so on, as well as examples like *TRAIN-RIDE* which involve Mere-Riding rather than controlled Riding.

These verb-noun relations can be included in the network, but we must first explain how generalisable lexical relations can be handled. Since a network is purely declarative, there is no place for procedures which create one lexeme on the basis of another - e.g. a rule which creates an adverb by taking an adjective and adding *-ly*. Instead, relations between lexemes must be handled by named functions such as "noun-of" or "adverb-of"; thus *QUICKLY* is the adverb-of *QUICK*. The generalisation in this case is that the adverb-of an adjective shares the adjective's stem but also has *-ly*. In the case of zero-derivation, the two lexemes simply share the same stem. For verb-noun pairs where the verb is considered basic, the noun-of the verb has the verb's stem - hence the fact that *WALK* can be used either as a verb or as a noun.

As far as meaning is concerned, the derived lexeme also has a sense which is based on that of its source lexeme. As already indicated, we assume that a concept may act as the sense of either a noun or a verb, so the noun-of the verb may have exactly the same sense as the verb; we believe this is true of many nominalisations (e.g. *ARRIVAL*, *LOVE*) but as the discussion above showed, it is not quite true of manner-of-motion verb-noun pairs like *WALK*, since the noun has the specific purpose of Life-Enhancing. This is the analysis presented in fig. 13, which can be paraphrased as follows: the sense of a Mom ("manner of motion") verb is a motion which has a manner, and its noun (i.e. its nominalisation) has the same stem and almost the same sense, but its sense has the added feature of aiming at Life-Enhancement.

---

[9] This explanation is supported by a minor dialect difference between the two authors regarding the possible meaning of the noun CYCLE: it can have either a concrete meaning (Bicycle) or an abstract one, as in "go for a cycle", but not both.

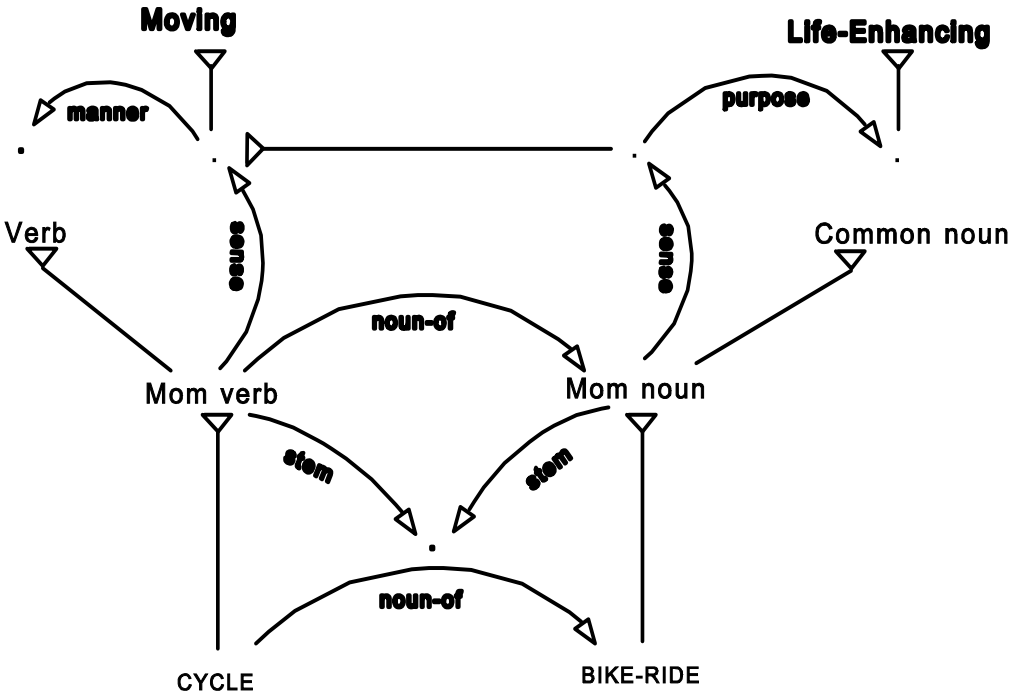

Figure 13

The other grammatical pattern that signals Life-Enhancing is illustrated by the following examples.

(27a) I went swimming with the kids.

(27b) Would you like to come flying?

(27c) We went mushroom-picking.

(27d) Please come paddling with me!

If you go swimming then you do it for fun or for fitness, and not to save your life. Therefore, you have not been swimming if you swim to the shore after being shipwrecked. Similarly for all the other examples; if you come flying, the destination is secondary, and most probably we will take off and land at the same place; if you go mushroom-picking, this is a pastime, not your way of earning a living; and so on.

Unlike the zero-derived nouns, manner-of-movement is only one of many possibilities, as illustrated by the example of mushroom-picking. Almost any activity which can be done for the sake of Life-Enhancing qualifies - fishing, painting, singing, etc. - but we cannot go (or come) working because Working conflicts fundamentally with the purpose of fun or fitness. However, there are some exceptions which we

cannot explain: we cannot go cooking, eating, coffee-drinking, chatting or sleeping, however life-enhancing these activities may be. Nor can we explain why some examples of the grammatical pattern define traditional activities which were (and are) done as jobs: we would go harvesting, hay-making or fruit-picking for money rather than for fun or fitness. However the main point is that we can easily go or come cycling, and that when we do so we are doing it for fun or fitness.

The grammatical pattern is more complicated than the zero-derived nouns discussed earlier. Indeed, we cannot claim to understand the grammatical pattern fully, as will become clear below. Its first part is one of the verbs *GO* and *COME*, which have their normal senses as deictically restricted verbs of direction-of-motion. You can either go or come swimming according to the deictic status of the destination, the swimming-place. However, in this pattern it is not possible to define the destination more precisely in the usual way:

(28a) I went (*swimming) to the seaside.

(28b) I went *(swimming) at the seaside.

The normal possibilities for spatial location seem to be replaced by those of the second word: *at the seaside* is possible in (28b) because it depends on *swimming*, but the expected dependency between *went* and *to the seaside* in (28a) is suppressed. We cannot explain how this verb can have a deictically specified direction but does not allow a normal direction adjunct.

The other part of the pattern is the "ing-word", which is much harder to analyse. This is partly because it is hard to distinguish from two other constructions: the non-movement pattern in (29a) below, which has negative overtones, and the pattern in (29b) where the ordinary *GO* or *COME* is combined with a participial adjunct.

(29a) Don't go/*come saying anything you'll regret.

(29b) He went/came/walked (to the woods) humming his favourite tune.

Another problem is that it is quite unclear what kind of word the ing-word is. Examples like (30a) suggest that it may be a participle, i.e. a verb, but (30b) shows that it can be a noun or adjective, because there is no verb *MUSHROOM-PICK*, as can be seen from (30c).

(30a) We went picking mushrooms.

(30b) We went mushroom-picking.

(30c) *We mushroom-picked.

Because of these uncertainties we shall not try to suggest a proper analysis for the syntax of the ing-word, nor for the semantics of *GO*/*COME*, so the diagram in fig. 14 is particularly provisional in syntax and in its treatment of the deictic meaning. In the semantics, "There" summarises the complex contrast between Going and Coming and is a reminder of the need to explain the ban on ordinary direction adjuncts that we saw in example (28a).

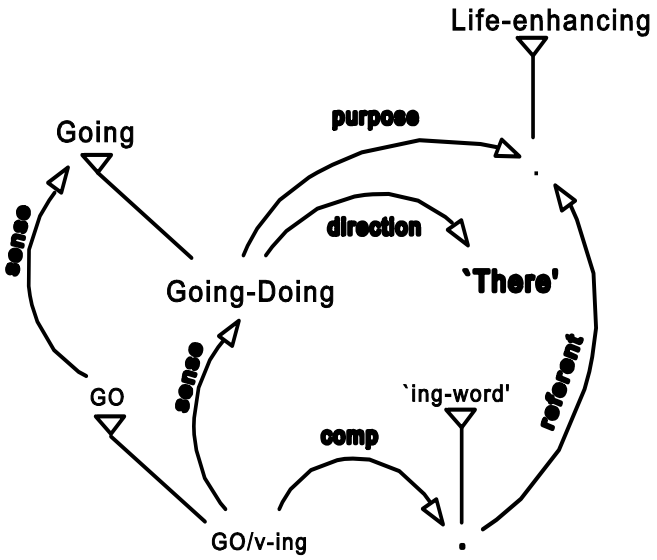

Figure 14

The details of the analysis are unimportant for our main point, which is the possibility of *go cycling* where the cycling has to be taken as something we do for fun or fitness. Like the zero-derivation pattern discussed earlier, this one defines the purpose as Life-enhancing, but as we have seen in (25) this is not the only purpose that can be forced on Cycling by the linguistic context. The flexibility of purpose suggests a semantic analysis of Cycling in which the purpose is left undefined, so that it can be filled in by the "construction meaning" as illustrated by the two constructions just discussed.

8. Conclusion

Although we have left some important threads dangling loose, we have been able to develop a reasonably insightful semantic analysis of Cycling which reveals its links to the notions Bicycle, Riding, Moving, Pedalling and a number of others for which there are no established names such as Life-enhancing and Mere-riding. The most

important gap in the semantic analysis is the concept Cycle (as in *the cycle of the seasons*), which we think is probably closer to Pedalling than to Cycling.

While discussing Cycling, however, we have tried to establish a number of general theoretical principles which we now summarise in order to stress that this kind of work cannot be done in a theoretical vacuum; indeed, we feel that it depends on a general theory of cognition.

- Concepts are recycled rather than duplicated - for example, the concept Bicycle is referred to directly in the definition of Cycling, rather than represented in that definition by some other concept which is similar to it.
- Recycling means that the sense of one word may be recycled as part of the definition of another; so the concept Bicycle which defines Cycling is the same concept as the one that is the sense of the lexeme *BICYCLE*.
- If non-language concepts are recycled as word senses, and vice versa, language contributes to the defining of concepts in the same way as other kinds of experience do, with its contribution ranging from zero (e.g. Bicycle) to almost totality (e.g. Riding or, perhaps, Life-enhancing).
- Network notations reflect recycling by assigning to each concept a single point on which many different relationships converge.
- Networks are a better model for conceptual structures than frames are because they do not imply boundaries.
- Networks and recycling are hard to reconcile with the traditional distinction between encyclopedic and dictionary meaning, because a network allows no such modular distinctions and any concept defined in the dictionary would also be recycled as a focus of encyclopedic relationships.

References

Anderson, John. 1993. “Problem Solving and Learning”. *American Psychologist* 48. 35-44.

Barsalou, Lawrence. 1992. “Frames, Concepts and Conceptual Fields”. In Adrienne Lehrer & Eva F. Kittay (ed.), *Frames, Fields and Contrasts. New Essays in Semantic and Lexical Organization*. Hillsdale: Lawrence Erlbaum. 21-74.

Fillmore, Charles. 1985. “Frames and the Semantics of Understanding”. *Quaderni di semantica* 6. 223-253.

Fillmore, Charles; Atkins, Sue. 1992. "Toward a Frame-Based Lexicon. The Semantics of RISK and its neighbours". In Adrienne Lehrer & Eva F. Kittay (ed.), *Frames, Fields and Contrasts. New Essays in Semantic and Lexical Organization*. Hillsdale: Lawrence Erlbaum. 75-102.

Hudson, Richard. 1985. "Some Basic Assumptions about Linguistic and Non-Linguistic Knowledge". *Quaderni di semantica* 6. 284-287.

Hudson, Richard. 1990. *English Word Grammar*. Oxford: Blackwell.

Hudson, Richard. 1995. *Word Meaning*. London: Routledge.

Hudson, Richard. 1998. *Encyclopedia of Word Grammar*. URL: http://www.phon.ucl.ac.uk/home/dick/papers.

Jackendoff, Ray. 1983. *Semantics and Cognition*. Cambridge: MIT Press.

Lakoff, George. 1987. *Women, Fire and Dangerous Things. What Categories Reveal about the Mind*. Chicago: University of Chicago Press.

Langacker, Ronald. 1987. *Foundations of Cognitive Grammar*, vol. 1. *Theoretical Prerequisites*. Stanford: Stanford University Press.

Lehrer, Adrienne; Kittay, Eva. 1992. "Introduction". In Adrienne Lehrer & Eva F. Kittay (ed.), *Frames, Fields and Contrasts. New Essays in Semantic and Lexical Organization*. Hillsdale: Lawrence Erlbaum. 1-18.

Levin, Beth. 1993. *English Verb Classes and Alternations. A Preliminary Investigation*. Chicago: University of Chicago Press.

Luger, George; Stubblefield, William. 1993. *Artificial Intelligence. Structures and Strategies for Complex Problem Solving*. Second edition. New York: Benjamin/Cummings.

Pinker, Steven. 1997. *How the Mind Works*. London: Penguin.

Reisberg, Daniel. 1997. *Cognition. Exploring the Science of the Mind*. New York: Norton.

Rosch, Eleanor. 1978. "Principles of Categorization". In Eleanor Rosch & Barbara B. Lloyd (ed.), *Cognition and Categorization*. Hillsdale: Lawrence Erlbaum. 27-48.

Schank, Roger; Rieger, Charles. 1974. "Inference and the Computer Understanding of Natural Language". *Artificial Intelligence* 5. 373-412.

Slobin, Dan. 1996. *Mind, Code and Text*. Unpublished manuscript, University of California, Berkeley.

Sowa, John. 1984. *Conceptual Structures. Information Processing in Mind and Machine*. Reading: Addison-Wesley.

Wierzbicka, Anna. 1985. *Lexicography and Conceptual Analysis*. Ann Arbor: Karoma.

Wierzbicka, Anna. 1988. *The Semantics of Grammar*. Amsterdam: John Benjamins.

Wierzbicka, Anna. 1996. *Semantics. Primes and Universals*. Oxford: Oxford University Press.