Metaphor in Discourse

Nicholas Asher
Department of Philosophy,
University of Texas,
Austin,
Texas 78712,
USA
nasher@bertie.la.utexas.edu

Alex Lascarides
Centre for Cognitive Science,
University of Edinburgh,
2, Buccleuch Place,
Edinburgh, EH8 9LW,
Scotland, UK
alex@cogsci.ed.ac.uk

October 17, 1997

1 Introduction

This paper focuses on metaphor and the interpretation of metaphor in a discourse setting. We propose constraints on their interpretation in terms of linguistic structures. Specifically, the constraints are based on a particular conception of the lexicon, where lexical entries have rich internal structure, and derivational processes or productivity between word senses are captured in a formal, systematic way (e.g., Copestake and Briscoe 1995, Pustejovsky 1995). By constraining metaphor in terms of these linguistic structures, we show that their interpretation is not purely a psychological association problem (cf. Lakoff and Johnson, 1980), or purely subjective (e.g., Davidson 1984). Recent accounts of metaphor within philosophy have not given systematic accounts of this sort (e.g., Black 1962, Hesse 1966, Searle 1979). We leave open the question of whether their insights are compatible with the theory proposed here.

Many have thought that the principles of metaphorical interpretation cannot be formally specified (e.g., Davidson, 1984). We'll attack this position with two claims. The first is that some aspects of metaphor are productive, and this productivity can be captured effectively by encoding generalisations that limit metaphorical interpretation in a constraint-based framework for

defining lexical semantics. This patterns in a loose sense with Lakoff and Johnson's (1980) view that the metaphor is productive. However, whereas they view the productivity as based in psychological concepts, we demonstrate that the productivity is essentially linguistic in nature. Indeed, from a methodological perspective, we would claim that the productive aspects of metaphor can give the linguist clues about how to represent semantic information in lexical entries.

Moreover, it is well known that domain knowledge and the way people structure fundamental concepts such as time and orientation influence metaphorical interpretation (Lakoff and Johnson 1980, and others). Our second claim takes this further, and we argue that *rhetorical relations*—such as *Elaboration*, *Contrast*, and *Parallel*, among others—that connect the meanings of segments of text together, can also influence the meaning of metaphor when they're presented in multi-sentence discourse, or even poetry. Through studying metaphor in discourse, we learn how to link lexical processing to discourse processing in a formal framework. We will give some preliminary accounts of how this link between words and discourse determine metaphor.

2 Metaphor in the Lexicon

First, we consider some examples of metaphoric productivity (cf. Moravscik, this volume). Lakoff and Johnson (1980) cite many examples of the phenomena, and argue that they are based on the conceptual system we live by. For example, they describe a CONDUIT metaphor, where ideas are treated as objects and language is viewed as a container:

(1) a. I gave you that idea.
b. It's difficult to put my ideas into words
c. When you have a good idea, try to capture it immediately in words.

But there are limits to the productivity, which must be captured. We must explain why (2a) are acceptable, but (2b) are odd, for example:

(2) a. Don't force your ideas into the student's head/onto the proletariat.

b. ?Don't force your ideas into the chairman/into the proletariat.

(1) and (2) provide evidence that phrases that denote linguistic expressions (e.g., *words*) and phrases that denote the cerebral parts of humans (e.g., *student's head*) can act as containers of ideas; hence the preposition *into* in these metaphors are acceptable (but *onto* would be odd). On the other hand, phrases that denote people (e.g., *chairman*, *proletariat*) cannot behave as containers (note that *onto the proletariat* is better than *into the proletariat* in (2). Why is this so?

There is also productivity in the way the class of *throw*-verbs combine with an abstract object such as abuse, as illustrated in (3a). But (3b) shows that this metaphorical phenomena is *semi*-productive.

(3) a. John threw/hurled/lobbed/flung/tossed abuse at Mary.

b. ??John batted/slapped/punted/tipped/flipped abuse at Mary.

An adequate theory of metaphor must describe how *hurl abuse* is semantically distinct from *throw abuse*, and also specify constraints on metaphorical interpretation that predict that (3a) are acceptable and (3b) aren't. Similar semi-productivity is illustrated in the class of *get*-verbs when combined with an abstract argument such as *new lease of life* in (4):

(4) a. John acquired/got/earned/gained/procured/stole/won a new lease of life.

b. ??John hired/rented/ordered/fetched/gathered a new lease of life.

A further source of productivity, and one which we will examine in detail in this paper, concerns the 'normal' vs. the metaphorical use of verbs describing change of location (CoL) in English and in French:

(5) a. Jean est entré en crise.
b. Jean est passé de l'autre côté de la loi.
c. He entered a blue funk/ came out of his blue funk.
d. Sam crossed the line of permissible behavior.
e. He deviated from the norm.
f. He skirted over the holes in the argument.
g. She stayed right on target.
h. You have now entered the Twilight Zone.
i. Jean est à la guerre.
j. He is outside (within) the law.
k. He is on top of the situation.
l. He was way off base.

(5a–h) employ verbs of motion or motion complexes, that in the typology described in Asher and Sablayrolles (1995) describe changes of location (CoL). (5i–l) employ complexes describing a position of the subject; these have a metaphorical behaviour very similar to the behaviour of CoL complexes, and so we'll loosely refer to these as CoL verbs too for the purposes of this paper.

These verbs conventionally apply to spatial locations. And yet the arguments to the verbs aren't of this type in (5). Following Lakoff and Johnson (1980), we don't treat the sentences in (5) as evidence that these verbs have a vague semantic sense, thereby increasing the burden on pragmatics to compute the specific senses in context, where one of these specific senses takes spatial locations as arguments. If we adopted this strategy, then we would lose a lot of information about how the conventional semantics of these verbs are constrained, and we would lose linguistic generalisations about the kinds of spatial arguments that the various CoL verbs can take. Rather, we assume that the CoL verbs in (5) are not being used in their conventional sense, but rather in some metaphorical one.

But there is a systematic relationship between the conventional sense of these verbs and their metaphorical senses. We agree with Lakoff and Johnson, that conventional vs. metaphorical ambiguity isn't homonymous, but rather there is a predictable relationship between the two senses. In particular, Asher and Sablayrolles (1995) describe a taxonomy for French CoL

verbs, which places restrictions on the kinds of movement that the CoL verb describes. For example, *entrer* takes the interior of some location or object l as an argument, and it describes an event of moving from some location near l to inside l.[1]

On the other hand, *arriver* is similar to *entrer* in that it takes the interior of some object or place l as an argument and describes an event of moving inside l, but it's also different in that it starts at a point *outside* the contextually defined "nearby zone" of l. In essence, the source of movement for *arriver* is 'further away' from the goal than it is for *entrer*, but the goals in each case are the same.[2] With this in mind, the sentences in (5) expose a lexical generalisation, which should be represented in a formal and computationally tractable manner. The lexical generalisation is this: the essential structure of the path described by CoL verbs—that is the type of change of location from a source to a goal via an intermediate path—is preserved in the metaphorical meaning. In other words, what distinguishes a particular CoL verb from other kinds of verbs in the conventional use, also distinguishes it from those verbs in their metaphorical use.

In particular, since the conventional use of *entrer* requires the interior of some physical location l as an argument, and it describes movement from near l to inside l, the metaphorical use of *entrer* requires an argument that needn't be some physical location, but it must have *extension*, so that we can talk about its interior. A bad mood or blue funk (mauvaise humeur) is such an example: it can be conceived of as having extension, because it is a state which extends in time. Hence (5c) is acceptable, and means that John 'transformed' from being in a good mood to being in a bad mood. If these predictions are correct, then just as conventional *entrer* is restricted in use (e.g., ?*John entered the line* is odd), its metaphorical interpretation will be restricted too. And this prediction is borne out, since it is possible to create incoherent examples. It is difficult to conceive of a line of permissible behaviour as having extension. The consequence is that you can cross it but

[1] Note that the words 'interior' and 'inside' are used loosely here, in the sense that l needn't be a container. Rather, any spatial zone that has some (at least two-dimensional) spatial extent has an interior or inside. So *entrer en France* or *entrer dans l'espace* are acceptable uses of *entrer* since both *France* and *l'espace* have interiors, and they are conventional uses since both *France* and *l'espace* are spatial zones.

[2] It's important to stress that the English *arrive* and *enter* have slightly different interpretations to their French counterparts. In particular, the difference between *arrive* and *enter* seems to depend in part on the functional or telic role of the spatial argument. We forego a detailed discussion of this here.

not enter it, as shown in (6):

(6) a. John crossed the line of permissible behaviour.
b. ??John entered the line of permissible behaviour.

Furthermore, it is difficult to conceive of a disturbing conclusion as having a "nearby zone", either spatial or otherwise. And this predicts that (7a) is acceptable, but (7b) is odd:

(7) a. We have arrived at a disturbing conclusion.
b. ??We have entered a disturbing conclusion.

Intuitively, the examples (3) and (4) follow these generalisations too, since these generalisations can explain the difference in acceptability between the a. and b. sentences. For example, what distinguishes *hire* from other *get*-verbs is that a contract is involved and the arrangement is temporary. According to our lexical generalisation, these things that distinguish *hire* from other *get*-verbs in the conventional sense also form the basis for distinguishing these verbs in their metaphorical senses. But it is difficult to conceive of a contract, metaphorically or otherwise, that underlies gaining a new lease of life.

In this paper, we will suggest how to exploit Asher and Sablayrolles (1995) classification of CoL verbs in French to explain these examples. We will encode the semantic information in a computational lexicon described within a typed feature structure framework, where other forms of productivity and polysemy have been studied extensively (e.g., Boguraev and Pustejovsky 1990, Briscoe *et al.* 1990, Copestake and Briscoe 1995, Pustejovsky 1991, 1995, Lascarides *et al.* 1996, and others). This framework is compatible with constraint-based grammars such as HPSG (Pollard and Sag 1994), and one can characterise certain types of productivity via lexical rules. We will propose a lexical rule called the Metaphor Lexical Rule, which encodes directly in the lexicon the restrictions on the metaphorical interpretation of words. Simply put, this rule will ensure that although the types of arguments of a verb can change in a metaphorical setting, the characteristics which differentiate it from other verbs in its class cannot. So, since *entrer* describes physical movement from nearby to inside a physical space in its conventional setting, it must describe some transition or movement from nearby to inside

something that needn't be a physical space (e.g., it could be a state or mood), in its metaphorical setting. Thus, the Metaphor Lexical Rule places *linguistic* constraints on the possible range of metaphorical interpretations of a particular word. This rule forms the basis for distinguishing between the acceptable and incoherent metaphors in (1–5).

Moravscik (this volume) argues that metaphorical senses are unspecific. This is supported in our theory by the fact that the Metaphor Lexical Rule won't always fully determine the interpretation of a metaphorical use of a word in a particular context; it simply limits the possibilities. However, the discourse context can influence the interpretation of metaphor, and in particular it can refine the unspecific senses that are determined by lexical information. We claim that some aspects of this influence of discourse on metaphor can be determined by computing how the metaphorical phrase is rhetorically connected to its discourse context. We will model this information flow between rhetorical relations and metaphor in a formal theory of discourse structure known as SDRT (Asher, 1993), which is distinctive in that it supplies an account of how rhetorical relations in discourse can affect semantic content of the propositions it connect. In some cases, this added semantic content will specify the nature of the metaphor.

3 The Semantics of CoL Verbs in French

Asher and Sablayrolles (1995) describe a taxonomy of French CoL verbs which groups them into ten classes. These classes are distinguished in terms of the seven zones that are depicted in Figure 1. The Z-inner-halo is typically defined to be the reference location provided by the complement to the verb (e.g., the object denoted in the PP); or this location may be contextually defined by the discourse context. The Z-outer-halo is space that is near to this location, and this is determined contextually, as is the 'far away' space Z-outer-most.

The ten classes stipulate which boundaries of these seven zones the movement described by the verb must cross, and the direction of this movement. These classes are depicted schematically in Figure 2. The boundaries of the zones that the movement described by the CoL verb must cross determines the kinds of arguments that the verb takes. All CoL verbs have a reference location l as one of their arguments, but their meaning differs in how this

reference location is exploited in constraints on the source, goal and intermediate path of the eventuality described. We have already described in words the restrictions on the movements described by *entrer* and *arriver*. The verb *sortir* resembles *entrer* in that it describes a motion involving l and a point in the near outside of l, but the movement goes in the opposite direction; l is the source of this movement and the point in the neighborhood is the goal.

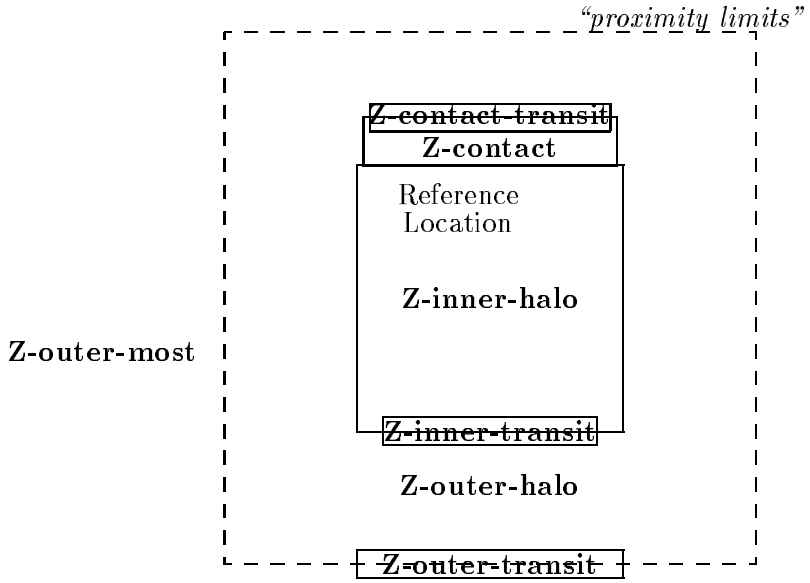

Figure 1: The Seven Generic Locations

The kind of information depicted in Figure 2 can be represented in a typed feature structure (TFS) framework for representing lexical semantic information. Through doing this, one can integrate the information about paths in the semantic component with other syntagmatic information provided in a constraint-based grammar such as HPSG (Pollard and Sag, 1994), so as to produce a compositional semantics for phrases and sentences involving CoL verbs. Using a constraint-based formalism also enables us to restrict the types of arguments that a verb can take, via the typing system that accompanies the TFSs.

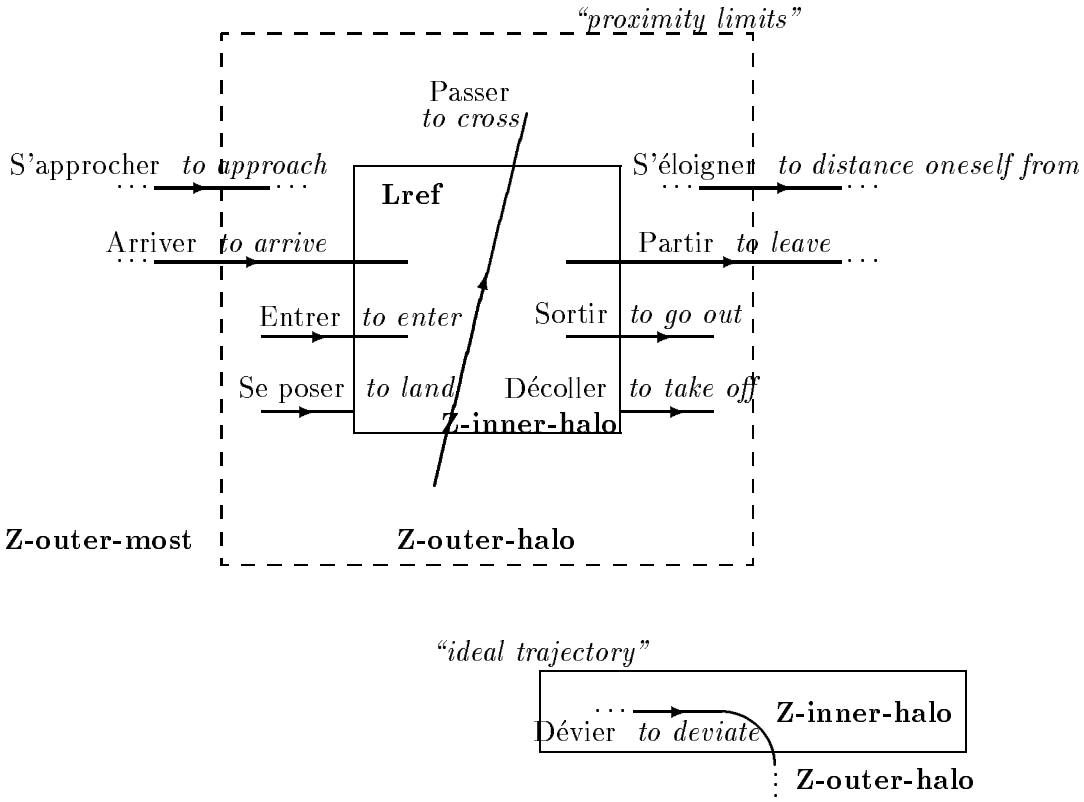

Figure 2: The 10 Groups of Verbs of **Change of Location**

HPSG is specially designed to capture the linking between syntactic structure and predicate argument structure in logical form. For example, the type **tverb** (standing for transitive verb) receives the type constraint given in Figure 3, and the type **dans-PPverb** (standing for a verb that takes a *dans*-PP) is given in Figure 4. We can exploit these generalisations in building lexical entries of verbs, since any verb of type **tverb** (e.g., *traverser*) or **dans-PPverb** (e.g., *entrer*) will inherit the linking given in Figures 3 and 4 respectively, and so we do not have to encode this information explicitly in each lexical entry. This is a model of lexical processing that has been suggested by many computational linguists (e.g., Copestake 1992) and in previous work, we have suggested how such a model can be linked to a

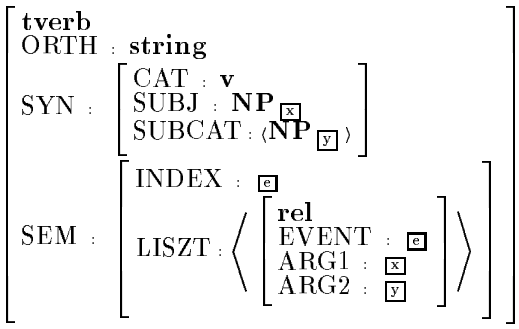

Figure 3: The Type Constraint for **tverb**

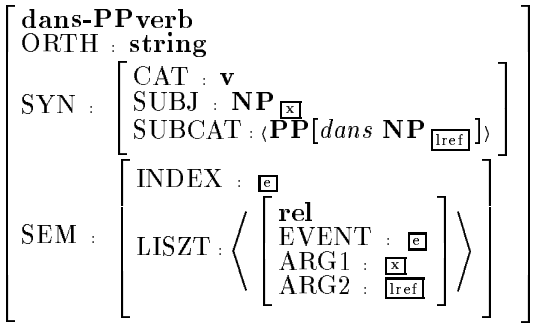

Figure 4: The Type Constraint for **dans-PPverb**

theory of discourse processing, to handle disambiguation (Asher and Lascarides 1995a, Asher and Sablayrolles 1995) metonymy (Lascarides 1995, Lascarides and Copestake 1997a) and zeugma (Lascarides, Copestake and Briscoe, 1996). We'll use it here to model metaphor.

Given the way we have classified CoL verbs, the semantics of each lexical entry will include information about the path of movement—from source through strict intermediate path (SIP) to goal—and various values which differentiate the different CoL verb classes. To capture these generalisations in a constraint-based formalism, we introduce more information about the semantics of the word than is standardly represented in HPSG, in a similar manner to that adopted in Copestake and Briscoe (1995) and Pustejovsky (1991, 1995). In particular, information about the type of event described by *aller* is given at the INDEX, via the *qualia structure*. A qualia structure for a given lexical entry provides information about the kind of object or event re-

ferred to by the word, along several dimensions. It reflects the fact that real world knowledge and word meaning are not necessarily distinct on this view of the lexicon (cf. Moravscik, 1990). Instead, lexical semantic information is a strictly limited fragment of world knowledge: this fragment interacts with knowledge of language and so is conventionalised in various ways. For example, a *qualia structure* for a word which denotes an object may include certain information about its appearance and its purpose. One dimension of description in the qualia structure is of special importance here. The FML dimension (standing for *formal*, and also known as *differentiae*) stipulates properties of the object or event being described, which distinguish it from other objects/events in a larger domain. We demarcate this information in the feature structure, so that we can define in the lexicon the generalisation that this information is preserved when the verb is interpreted metaphorically, as we'll shortly see. It should be stressed, however, that the value of the FML feature may not include *all* real world knowledge that's relevant to distinguishing objects or events from each other. Information in the *qualia*, although closely connected with real world knowledge, does not subsume it. It includes only a limited subset, for which there must be syntagmatic evidence that the information is part of the conventional meaning.

In the case of CoL verbs, the differentiae describe the properties of the source, intermediate path and goal of the movement described by the verb. Indeed, the fact that these properties constrain the syntax of their conventional use (and their metaphorical use), as shown in (6), provides evidence that they're part of the conventional meaning of the verbs. To encapsulate these properties of movement in the lexical entries for verbs, we first place constraints on the type **CoL-verb**, which ensure that each **CoL-verb** has features that define the source, SIP and goal of the path, and ensure that the object that is defined to be at the source, SIP and goal of the movement is the thing denoted by the first argument in the semantic structure; which in turn is the subject of the sentence given the linking provided by the type constraints on verbs discussed above. This is given in figure 5.

Second, we must ensure that the spatial zones, which will ultimately be introduced as arguments to the source, SIP and goal of the path in any given lexical entry for a CoL verb, are defined by the complement when it exists (for note that the value of SEM:LISZT:ARG2 may be contextually defined). We represent these as implicative constraints. In words, (8) stipulates: if a TFS has some path π with value **Z-inner-halo**, then this value is computed

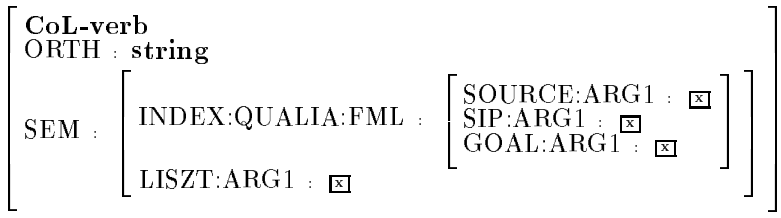

Figure 5: Type Constraints for **CoL-verb**

as a (contextually defined) function f_{ih} of the value **l**, which is the second argument in the predicate argument structure in logical form. Similarly for **Z-outer-halo** and **Z-outer-most**.

(8) a. $\langle\pi\rangle$ = **Z-inner-halo** $\rightarrow$
(**Z-inner-halo** = f_{ih}(**l**)$\wedge$
$\langle$SEM:LISZT:ARG2$\rangle$**l**)

b. $\langle\pi\rangle$ = **Z-inner-halo** $\rightarrow$
(**Z-outer-halo** = f_{oh}(**l**)$\wedge$
$\langle$SEM:LISZT:ARG2$\rangle$**l**)

c. $\langle\pi\rangle$ = **Z-outer-most** $\rightarrow$
(**Z-inner-halo** = f_{om}(**l**)$\wedge$
$\langle$SEM:LISZT:ARG2$\rangle$**l**)

Finally, we must encode the fact that when a CoL verb is used conventionally, the first argument in the semantic structure must be a mobile object, and the second one must be a physical space, so that we guarantee that the CoL verb describes an event where a mobile object moves in physical space in some way. We assume that this selectional restriction on the arguments to the conventional sense of a CoL verb are encoded via type constraints on the value of SEM:INDEX for this verb. So we assume that, at least in their conventional setting, CoL verbs invoke an event e in semantics which is of type **CoL-event**, and this type has constraints on it, which restrict the types of the other arguments in the semantic structure: the first argument must be a **mobile-object** and the second one must be a **physical-space**. This is specified in the implicative constraint (9):

(9) $\langle$SEM:INDEX$\rangle$**CoL-event** $\rightarrow$
($\langle$SEM:LISZT:ARG1$\rangle$**mobile-object**$\wedge$
$\langle$SEM:LISZT:ARG2$\rangle$**physical-space**)

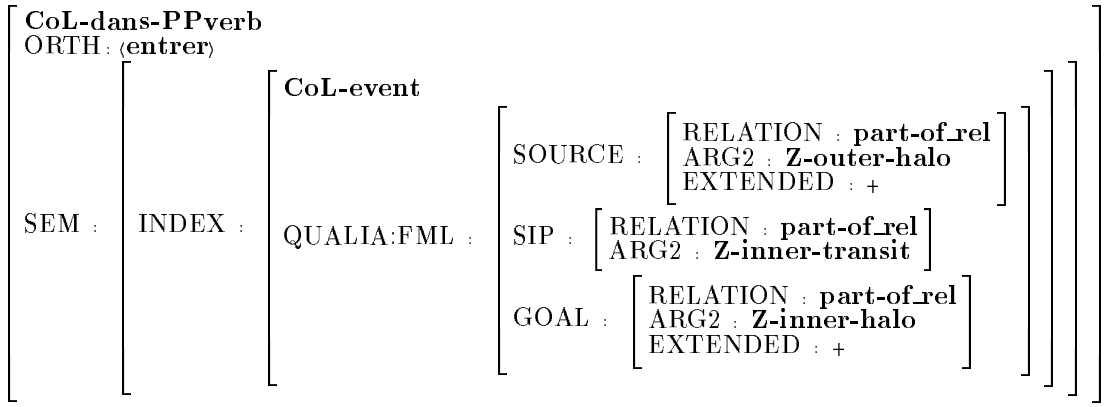

Figure 6: The TFS for *entrer*

We encode these selectional restrictions on this 'semantic' type **CoL-event**, rather than via the type constraints on **CoL-verb**, for two reasons. First, many have argued that selectional restrictions must be semantically rather than syntactically determined. Second, as we have shown, **CoL-verb**s need not be used conventionally; they can be used metaphorically. And in such cases, the syntax remains the same but the types of the arguments in the semantic structure change. We can represent this preservation of syntax but change in selectional restrictions if we keep constraints on syntax and selectional restrictions separate in this way.

For example, the CoL verb *entrer* is a **CoL-dans-PPverb**, which is a subtype of **CoL-verb** and **dans-PPverb**. Consequently, *entrer* inherits the constraints from both these types. In particular, it inherits the linking constraints on **dans-PPverb** provided in Figure 4, and the constraints given in Figure 5. So the specification of the lexical entry for the conventional use *entrer* is given in figure 6.

Note that in this TFS, the event is typed to be **CoL-event**. Therefore, the subject and object are restricted to be of types **mobile-object** and **physical-space** respectively via the implicative constraint (9). Taking this and the type constraints on **dans-PPverb** and **CoL-verb** into account, the expanded version of the lexical entry for (the conventional use of) *entrer* is given in Figure 7. Note that by (8), the values for **Z-inner-halo** and **Z-outer-halo** in this TFS must be contextually defined via the reference location given by the complement, and that these values must be extended.

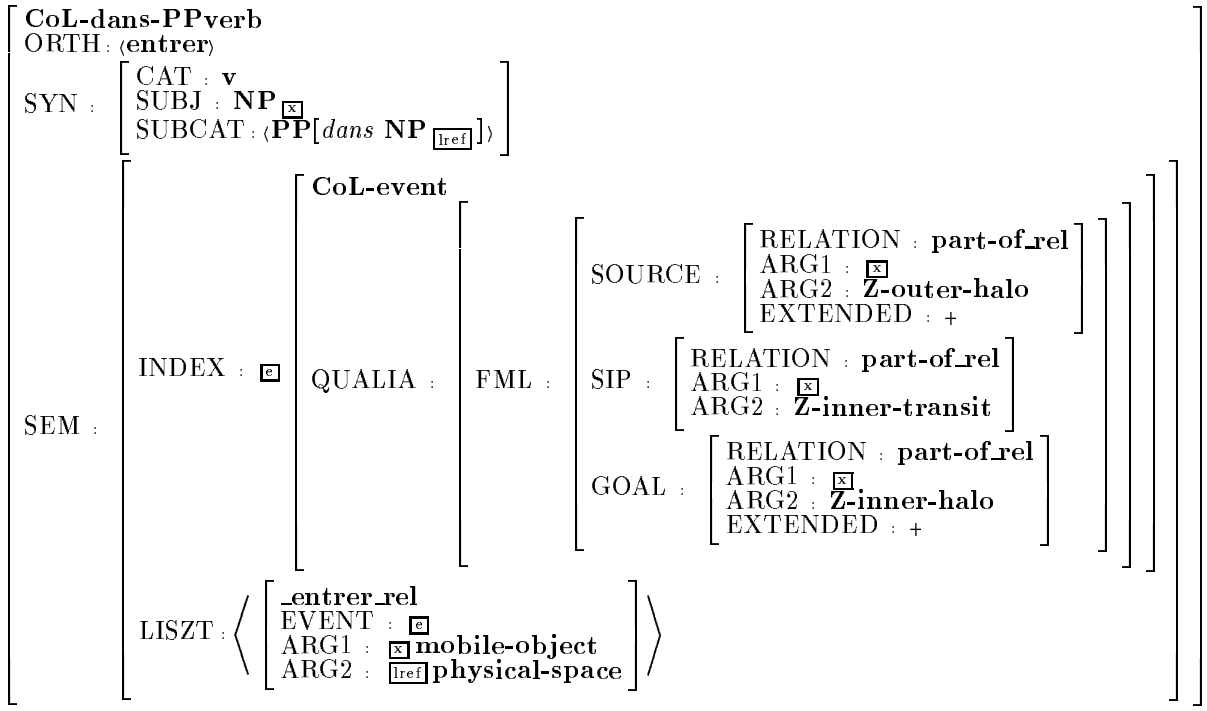

Figure 7: The Expanded Feature Structure, derived from the conventional lexical entry for *entrer*

In contrast, the **Z-inner-transit** need not be extended. This reflects the fact that this zone need not be contextually defined for *entrer* to be used coherently. However, if it can be defined (e.g., as the door of the room, for *Jean est entrée dans la salle*, or customs for *Jean est entrée en France*), then the lexical structure of *entrer* predicts that Jean entered through this zone—i.e., through the door or through customs.[3]

So the verb *entrer* (and other verbs grouped with *entrer*) describe a motion of an object [x], which by the semantic constraints on **CoL-event** must be a **mobile-object**, and this object goes from a place outside of, but near to, the inside of a location of reference (given by the complement PP if it

[3]This aspect of the meaning of *Jean est entrée dans la salle* can be overridden in sufficiently rich discourse contexts, and so it provides a good candidate for the use of lexical defaults, in the style of Lascarides and Copestake (1997a, 1997b). But we forego a discussion of that here, since it will divert us from the main purpose of this paper.

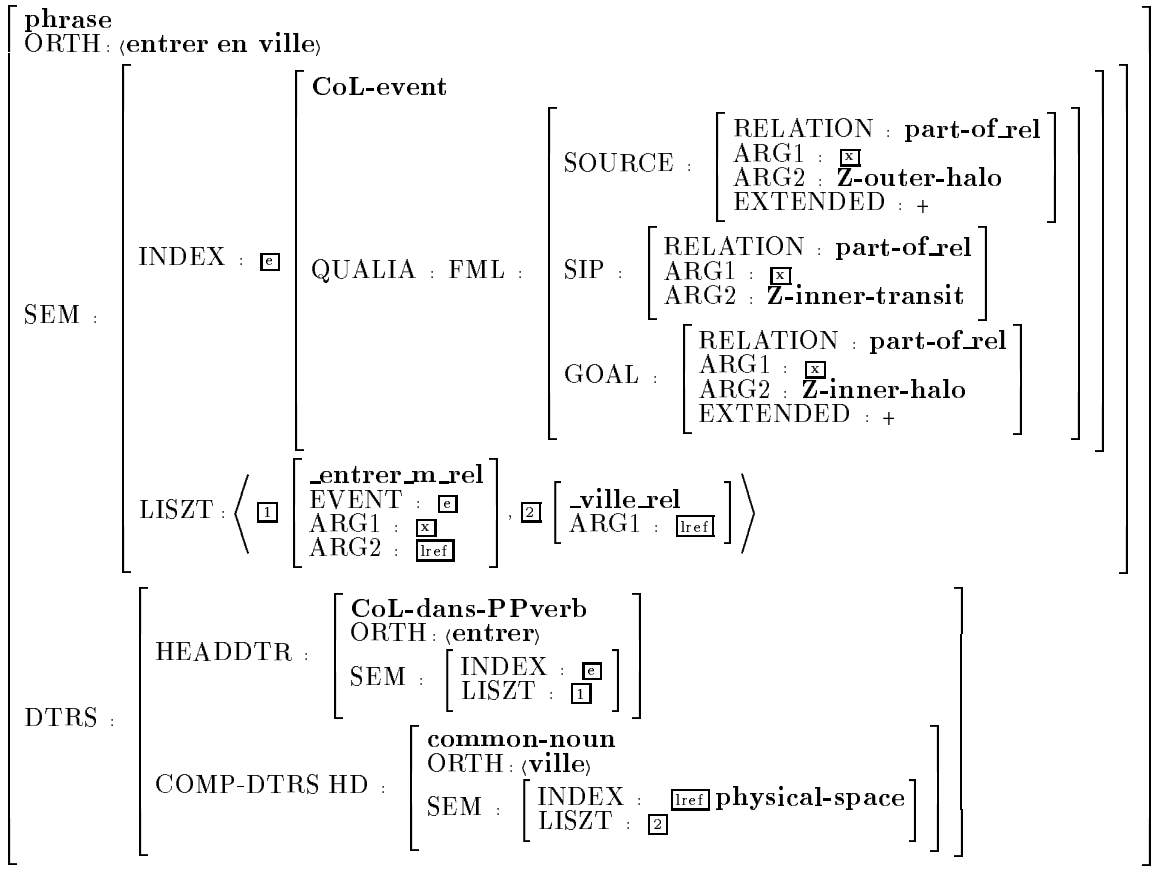

Figure 8: The TFS for *entrer en ville*

exists), where this location must be **physical-space**, to the inside of that location of reference, via an inner-transit zone, if it can be defined. So, the TFS for the phrase *entrer en ville* is given in Figure 8. For the sake of brevity, we have omitted the syntactic component. For the purposes of this paper, we also treat *en* as indistinguishable from *dans le*.

4 Lexical Restrictions on Metaphor

All lexical entries for the conventional sense of CoL verbs are encoded so that the type of event on SEM:INDEX is **CoL-event**. Consequently, by the

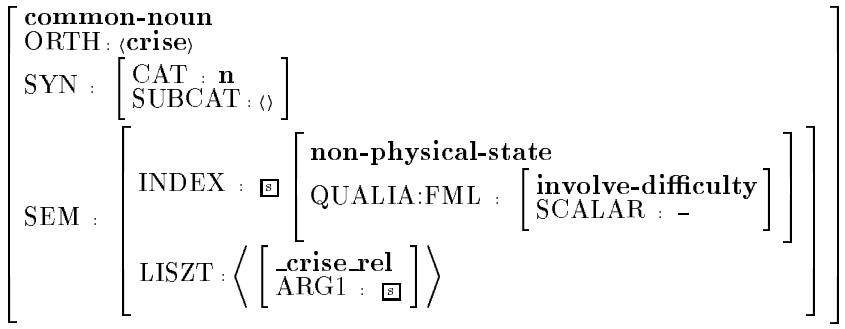

Figure 9: The Feature Structure for *crise*

type constraints given in (9), lref must be a physical space. The argument lref may be picked up from a syntactic complement in the sentence such as a PP, or it may be specified through contextual information in the preceding discourse. It is ultimately used, via the implicative constraints in (8), to define the properties of the movement described by the CoL verb, which appear in the SEM:QUALIA:FML component of the lexical entry.

Because the argument lref is restricted to be physical-space, all the lexical entries for the conventional sense of CoL verbs cannot be the ones that are used to represent the sentences in (5). Attempting to compose the TFSs in these cases within the grammar will fail, because the relevant argument that fills the lref slot doesn't have the appropriate type on the path SEM:INDEX in its lexical entry. In particular if we consider the example (5a), where the argument which fills in lref is *crise* or *crisis*, there is a type clash triggered by the type constraints on **CoL-event**: this requires the semantic index of lref to be **physical-space**, but filling lref with the PP complement *crise* makes it **non-physical-state**, as shown in the lexical entry for *crise* given in Figure 9. In words, this TFS defines *crise* to be a non-physical state, and what differentiates it from other non-physical states in its subclasses and superclasses is that it's non-scalar and involves difficulty.

As we've mentioned, the sentences in (3), (4) and (5) expose a lexical generalisation: what distinguishes a word from other kinds of words in the conventional use, also distinguishes it from those words in their metaphorical use. So, in particular, the metaphorical senses of CoL verbs preserve the properties of the path which describes the motion, from the source to a goal via an intermediate path.

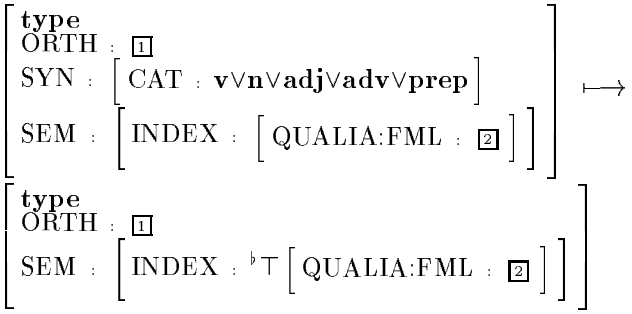

Figure 10: The Metaphor Lexical Rule

Ideally, we should capture this constraint as a generalisation in the lexicon: the lexicon should predict which aspects of a lexical entry can vary in metaphorical interpretation, and which can't. We have demarcated the semantic properties of words that distinguish it from other words in a larger domain in the feature SEM:INDEX:QUALIA:FML (cf. Pustejovsky 1991, 1995). Because of this, we can capture the constraint that the value on this path is preserved in a metaphorical setting. We specify this constraint in a *lexical rule*, which we represent as a function between TFSs in the spirit of Briscoe and Copestake (1992). The relevant rule is given in figure 10.

The Metaphor Lexical Rule constrains the metaphorical senses of words. The left hand side (LHS) TFS constrains which words can have metaphorical senses, and the RHS TFS constrains the metaphorical senses in the following ways: orthography, syntax, the linking constraints given by the root type, and the differentiae for the conventional senses of words are preserved in the metaphorical senses.[4] However, regardless of the (semantic) type of the index of the conventional sense of the word, it can be anything in the metaphorical sense. This latter piece of information is captured by the value $^{\flat}\top$ on the path SEM:INDEX in the RHS TFS. In line with Mereurs (1995), we use the notation $^{\flat}\top$ to indicate that the type $\top$ doesn't unify with the input

[4]This account assumes that all the information in the differentiae is preserved in the metaphorical meaning. This may seem like a strong claim. However, as we've mentioned, the differentiae don't include all of the real world knowledge that people recruit to distinguish between objects and events; it only includes that information for which there is evidence that it affects linguistic structures. So preserving all aspects of the differentiae is a weaker restriction than one might expect.

TFS to the lexical rule, but rather $\top$ *replaces* it.

From a formal perspective, these properties arise from the way inputs and outputs to the rule are computed. One uses unification and default unification respectively (e.g., Briscoe and Copestake, 1997). First, a TFS can be input to the rule only if it is unifiable with the LHS TFS in that rule. So, because of the value given on SYN:CAT in the LHS TFS, the Metaphor Lexical Rule will take as input only words which are verbs, nouns, adjectives, prepositions, and adverbials. But it predicts that quantifiers (e.g., *all*), auxilliaries (e.g., *do*) and connectives (e.g., *but*, *and*) don't have metaphorical senses.

Default unification implements the requirement that information in the input should carry over to the output of the lexical rule, provided that it is consistent with the output type defined by the rule, and any further constraints specified in the rule. Reentrancies as well as matching types in the lexical rule across the LHS and RHS TFSs capture equivalences between the conventional and metaphorical senses.[5] To compute the output of the rule, given the input, one replaces the reentrancies and matching types (e.g., the root type **type** in the Metaphor Lexical Rule) in the RHS TFS with the values on these paths given by the input TFS. This is then treated as indefeasible information. So, for example, if the input to the rule is the lexical entry *entrer* given in Figure 6, then the indefeasible information that's used to compute the output to the rule has root type **CoL-dans-PPverb** (and so the output will inherit the linking constraints associated with verbs of this type); the value $\langle$**entrer**$\rangle$ on ORTH; the same value on SEM:INDEX:QUALIA:FML as that given in Figure 6, and the value $^{\flat}\top$ on SEM:INDEX (the $\flat$ sign indicating that $\top$ will replace any value given on this path by the input TFS, rather than unify with it, as we mentioned before). The input TFS is then taken to be persistent default information, and the output of the rule is computed by default unifying these things together, as described in Lascarides and Copestake (1997b) (cf. the definition of lexical rules given in Briscoe and

[5]Note that in line with Copestake and Briscoe (1995), this rule can be glossed as a single feature structure typed **metaphor-lexical-rule**, where the TFS on the LHS of $\longmapsto$ is the value of the feature INPUT or 1, and the TFS on the RHS is the value of the feature OUTPUT or 0. Thus the reentrancies in this rule can be assigned the standard semantics. Treating lexical rules as TFSs in this way allows us to specify a subtype of this lexical rule, which could restrict the metaphorical interpretation of subclasses of the class of words that can be interpreted metaphorically even further. Such rules would inherit from the TFS typed **metaphor-lexical-rule** the constraint that the differentiae must be preserved. But these subtype of the rule could specify further constraints in addition to this one.

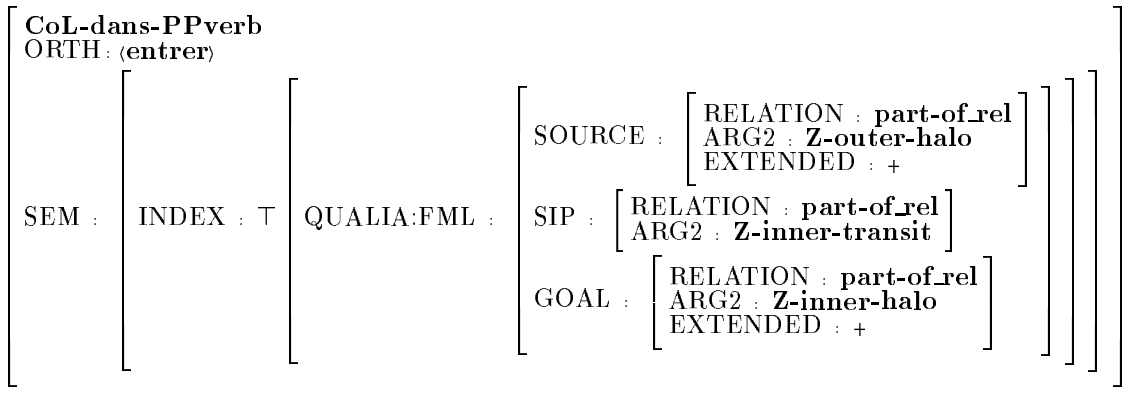

Figure 11: The output of the Metaphor Lexical Rule for *entrer*

Copestake, 1997). We take the default information to be persistent, so that any qualia information which survives in the metaphorical sense because it was consistent with the indefeasible information (e.g., the value of the differentaie) remains marked as default. As shown in Lascarides and Copestake (1997b), it could therefore subsequently be overridden, when it is combined with incompatible (indefeasible) qualia information from other lexical entries in the phrase. However, because the differentiae in the metaphorical senses of words contain indefeasible information, this information cannot be overridden by incompatible information in other lexical entries in the phrase.

For example, the output of the Metaphor Lexical rule for the TFS in Figure 6 is given in Figure 11. Thus the metaphorical entry for *entrer* is the same as the conventional entry, save that **CoL-event** is replaced with $\top$ (note that $\top$ replaces **CoL-event** rather than default unifies with it, as explained earlier). Therefore, the type constraints on **CoL-dans-PPverb** apply to the metaphorical sense of *entrer*. In contrast, the type constraints on **CoL-event**, which applied to the conventional sense, do not apply to the metaphorical sense, because this type has been replaced with $\top$. Thus the TFS representing the metaphorical entry for *entrer*, once it is expanded out via inheritance with the relevant type constraints, looks exactly like the entry given in Figure 7, save that: **CoL-event** is replaced with **event** (it must be **event** because of the type constraints on **dans-PPverb**); **mobile-object** and **physical-space** no longer feature at all; and **_entrer_rel** is replaced with **rel**. Moreover, were we to add any other features of *entrer* to the conventional entry in Figure 6, such as a telic role under SEM:INDEX:QUALIA,

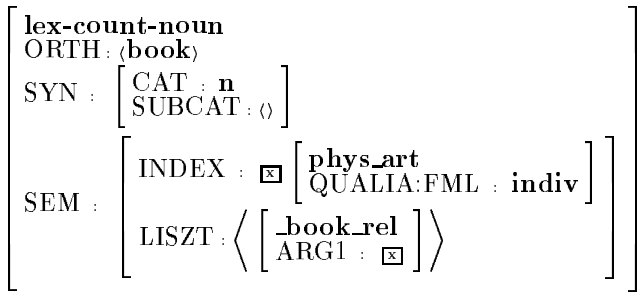

Figure 12: The Typed Feature Structure for *book*

then these would survive as persistent defaults in the metaphorical entry. That is, they would appear by default and marked as default in the TFS for the metaphorical sense of *entrer*.

The consequence of changing the type on SEM:INDEX is that the denotation of the metaphorical sense of the word is potentially of a different kind from the denotation of its conventional sense. Moreover, as we have mentioned in the case of *entrer*, any constraints produced by this type in the conventional entry don't survive in the metaphorical entry. For CoL verbs in general, this means that their metaphorical senses can have subjects and complements which aren't mobile objects or physical spaces respectively. Similarly, given the lexical entry for a noun like *book* given in Copestake and Briscoe (1995), as shown in Figure 12, *book* is still a lexical count noun when it's used metaphorically—and so has the syntax-semantics linking associated with this type—and it still refers to an individuated object (since this is the value on the FML path), but it need not be a physical artefact, since the type on SEM:INDEX changes from **phys_art** to $\top$.

Note also, that since the Metaphor Lexical Rule will ensure that any qualia information apart from the value on FML survives in the metaphorical entry only by default, the kinds of words that metaphorical senses can combine with is different (and in general a larger set) than the ones that the conventional sense can combine with. For example, QUALIA:CST defines the stuff that an object is made up of (Pustejovsky 1991, 1995). The Metaphor Lexical Rule ensures that the value of this on the metaphorical sense of a word is marked as default. So consider *cloud*, for example. In its conventional sense, the value on QUALIA:CST is marked to be of type **physical-cum** (so clouds are made up of a physical cumulation of stuff; cf. Copestake and Briscoe

1995, Lascarides *et al.* 1996). So this value survives in the metaphorical entry for *cloud*. But it is marked as default, and can be overridden by incompatible information in other lexical entries in the phrase, or by pragmatic information in the discourse context (Lascarides and Copestake, 1997a). This happens in *cloud of suspicion*: this refers to something made up of (non-physical) suspicion, rather than something physical.

It should be stressed that the Metaphor Lexical Rule will not determine the interpretation of a metaphorical use of a word on its own. This is because, as is standard in constraint-based approaches to lexical semantics, several of the values in the representation of the lexical entry are underspecified in some sense. In particular, the value on SEM:INDEX is $\top$—the most general type of all. We have utilised this underspecification in order to capture the fact that the metaphorical uses of words are in general less determinate than their conventional counterparts, about the semantic types of words that they can combine with to form phrases (cf. Moravscik, this volume). However, it's not completely unconstrained. The Metaphor Lexical Rule captures linguistic restrictions: orthography, linking and the differentiae must all be preserved. This constrains the possible senses of metaphorical uses, and pragmatic information must be used to compute the preferred interpretation in the given discourse context from this range of possibilities. So in general, one cannot compute the metaphorical interpretation of a word in context in the absence of open-ended, nonmonotonic pragmatic reasoning. We'll examine how pragmatic information aids the interpretation of metaphor in section 5. At any rate, the Metaphor Lexical Rule serves to restrict the possibilities *monotonically* in the lexicon.

In the case of CoL verbs, the Metaphor Lexical Rule will capture the fact that the distinctions among the CoL verbs and verb complexes, described in Asher and Sablayrolles (1995), are preserved in the metaphorical sense extensions. As we've seen, we assume the *differentiae* of CoL verbs stipulate the constraints on the source, goal and intermediate path that distinguishes it from other verbs. And the Metaphor Lexical Rule ensures these constraints are maintained in the metaphorical senses, but the types of the arguments the verbs take won't be constrained to be **mobile-object** and **physical-space** anymore.

As an illustrative example, let's examine how the Metaphor Lexical Rule contributes to the analysis of (5a) in the monotonic component of the grammar. As we've already mentioned, the TFSs for the conventional senses of

entrer and *crise* given in figures 7 and 9 cannot be unified. Indeed, the same holds for all the examples we have seen in (5). However, we can use the Metaphor Lexical Rule to produce TFSs that will unify. There are three alternatives. First, we can create a metaphorical sense for *entrer* via the Metaphor Lexical Rule which will unify successfully with the TFS for *crise* given in Figure 9. Second, we could allow *crise* to be input to the Metaphor Lexical Rule rather than *entrer*; or third, we allow both *entrer* and *crise* to be input. Allowing both to be input would result in more work for the pragmatic component, because more information would be removed from the TFSs by the Metaphor Lexical Rule. On the other hand, preserving the TFS for *entrer* and allowing *crise* to undergo the mapping defined by the Metaphor Lexical Rule would mean that *crise* would have to be interpreted as a **physical-space** (since the conventional sense of *entrer* would constrain it to this), which is differentiated from other physical spaces in that it involves difficulties (because the output of the Metaphor Lexical Rule on *crise* would preserve the differentiae). This interpretation of *crise* as a physical space would have to be pragmatically determined, and it is unclear how one could conceive of a physical space involving difficulty. Therefore, we will examine in detail here the analysis produced by allowing *entrer* to undergo the mapping, and *crise* to stay as it is. With this assumption, and using the TFSs given in Figures 11 and 9, the TFS for the phrase *entrer en crise* is given in figure 13 (for the sake of simplicity, we have again omitted syntactic information).

But there is more work to be done. The arguments to CoL verbs in their metaphorical use must be interpretable as locations in some qualitative space, so that the zones Z-outer-halo, Z-inner-halo and possible Z-inner-transit, which feature in the *differentiae* of the verb, can be computed in line with the implicative constraints given in (8). This computation will determine the (non-spatial) path given by the metaphorical phrase. That is, we must (contextually) compute the zones need in the differentaie from *crise*, so that we can compute the source, intermediate path and goal of the (non-spatial) path.

There are two strategies we could adopt at this point. The first is to define lexical rules that produce sense extensions of nouns that allow them to be interpreted as locations in some qualitative space. Indeed, these lexical rules could be subtypes of the Metaphor Lexical Rule and inherit from that rule the fact that the value on QUALIA:FML must be preserved. The second

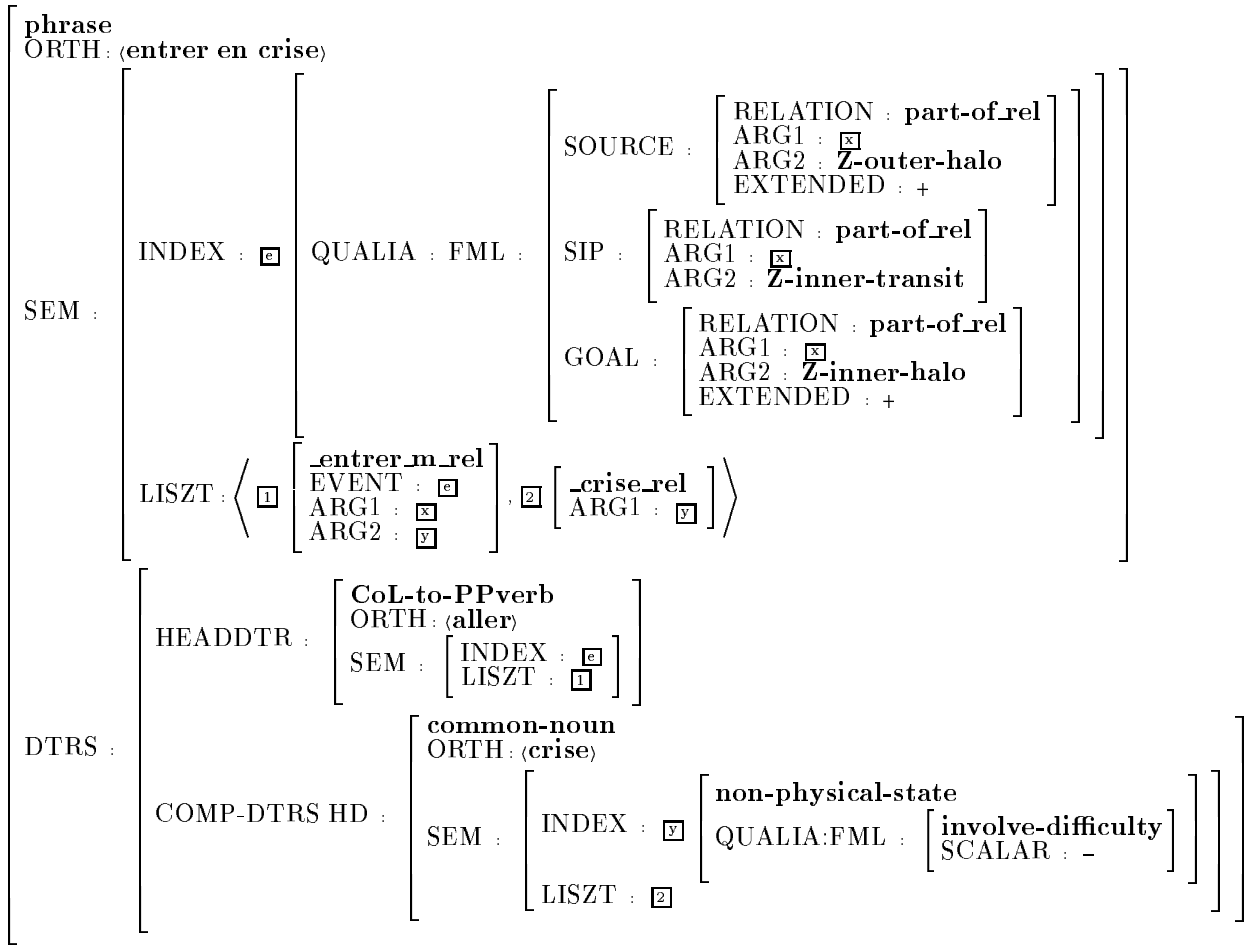

Figure 13: The TFS for the phrase *entrer en crise*

is to assume that these locations are contextually determined, rather than assuming that they are computed in the lexicon itself. We adopt the latter strategy for two main reasons. First, the lexical rule strategy would proliferate senses of nouns in an unnecessary manner. There is no *prima facie* reason for assuming that *crise* as it is used in (5a) is a different sense from *crise* as it is used in (10), for example.

(10) La crise a commencé en Septembre 1939.

Second, even outer most zones and nearby neighbourhoods for spatial locations such as *l'école* are contextually determined, and so assuming that the zones are contextually determined for non-spatial nouns as well produces a more uniform analysis. Therefore, we need to formally represent generalizations in the pragmatic component, which enable one to compute the necessary zones of qualitative space for non-spatial nouns.

We are on less certain ground when interpreting nouns as locations in qualitative space, since we have not made an exhaustive study of them. However, we tentatively propose two generalizations. First, we claim that any non-scalar common noun or NP α that has an antonym β allows us to construct locations in qualitative space as follows: the Z-inner-halo is α itself, and the Z-outer-halo is the antonym. This rule is formally defined in `Zones for Non-Scalar Nouns`:

- `Zones for Non-Scalar Nouns:`
  $(noun(\alpha) \wedge antonym(\alpha, \beta)) \rightarrow (Z\text{-}inner\text{-}halo(\alpha, \alpha) \wedge Z\text{-}outer\text{-}halo(\alpha, \beta))$

It is important to stress that this rule doesn't enable us to compute outermost zones (or Z-outer-most) for non-scalar nouns. And if we can't determine what the zones are, we must assume they have no extension; they can be thought of as absent.

The second generalization for computing zones involves scalar nouns: for example *bad mood*. If α is part of a *scalar* system (e.g. good, bad, better, worse), then we claim that the outermost zone is the antonym, and the *near* neighbourhoods of the qualitative space can be constructed via the α-like objects γ in the scale. This rule is given in `Zones for Scalar Nouns`:

- `Zones for Scalar Nouns:`

$$(noun(\alpha) \wedge antonym(\alpha, \beta) \wedge scalar(\alpha) \wedge scale(\alpha, \gamma)) \rightarrow$$
$$(Z\text{-}inner\text{-}halo(\alpha, \alpha) \wedge Z\text{-}outer\text{-}halo(\alpha, \gamma) \wedge Z\text{-}outer\text{-}most(\alpha, \beta))$$

So, for example, for the NP *bad mood*, the Z-inner-halo is the bad mood, the Z-outer-most is a good mood, and the near neighbourhoods γ are moods that go progressively from good to bad, as one goes from the Z-outer-most towards the Z-inner-halo.

These two rules enable us to compute zones for the verb's arguments, when the lexical entry of the verb produces a logical form which requires such an interpretation. However, it should be stressed that we do not consider these generalisations to form part of the lexicon. Rather, these rules are viewed as pragmatic information on how one can compute locations in qualitative space from nouns, as and when they're required for interpreting a phrase.

As an example of how these generalizations contribute to the metaphorical interpretation of CoL phrases, consider example (5a) again. *Crise* is a non-scalar noun with an antonym, and therefore the rule `Zones for Non-scalar Nouns` applies, whereas `Zones for Scalar Nouns` does not. Thus, we can infer that the Z-inner-halo is a state of crisis, the Z-outer-halo is a state of 'equilibrium', but the outermost zones aren't computable at all. Luckily, the FML feature on *entrer* doesn't require this zone to have extent, and so there is no anomaly. Thus we predict that (5a) is acceptable, and it means that Jean was first in a state of equilibrium, and then at some point after he was in a state of crisis. Similarly, *entrer dans une nouvelle crise* and *sortir de la crise* are also acceptable.

Moreover, this theory predicts that any CoL verb which requires that the source or goal be a Z-outermost cannot combine with a non-scalar noun such as *crise* to produce an acceptable metaphor. This is supported by the linguistic data. For example, the CoL verb *aller* requires the source of the movement to be the **Z-outermost**, as shown in figure 2. So the metaphorical interpretation of *aller* is lexically restricted by the Metaphor Lexical Rule to require the source to be this region. Thus we predict, correctly, that *aller en crise* and *aller à une crise* are both odd, contrasted with *aller en ville* and *aller à la ville*, which are both acceptable. This is because the source of the metaphorical movement, which is lexically constrained to be **Z-outermost**, cannot be computed from *crise*.

The Metaphor Lexical Rule combined with our observations about the ontological status of certain locations also predicts that (5d) is fine, while (11)

is not.

(5) d. Sam crossed the line of permissible behaviour.

(11) *Sam has entered the line of permissible behavior.

A line, even one in a qualitative space defined by possible behaviors (a scalar concept again), is something that can be crossed but that cannot be entered, since lines in our commonsense geometry have no extension (for a discussion see Aurnague and Vieu 1993). Thus, the lexicon will fail to provide TFSs that can combine in the monotonic component of the grammar to produce an acceptable representation of (11).

By studying certain aspects of metaphorical productivity, we have learned about the kinds of features that we need in the semantics: representing *differentiae* would be useful for capturing lexical constraints on the degree of variation permitted in metaphorical interpretation.

As a further example, consider *hit*-verbs such as *hammer*. The metaphorical sentences (12a) are acceptable, whereas (12b) is odd:

(12) a. John hit/hammered/flattened/crushed Mary with abuse.
b. *John tapped Mary with abuse.

Our lexical restriction on metaphor can predict this. The conventional sense of *tap* is distinct from *hit*, *hammer*, *flatten* and *crush*, in that it describes an action which (by default) does not affect the essential characteristics—such as shape—of the patient. This difference between *tap* and other verbs in its class is represented in its SEM:INDEX:QUALIA:FML value. Assuming that *abuse* is coherently used only if the patient is affected in some essential way, the Metaphorical Lexical Rule predicts the difference in acceptability between (12a), where such an affect can be computed via the differentiae of the metaphorical senses of the verb (which are the same as in their conventional senses), and (12b) where no such affect can be computed because the metaphorical sense of *tap* fails to provide it.

A further example of metaphorical productivity comes from Ortony (1979): adjectives which apply to physical objects, can be applied in a metaphorical sense to humans: *straight*, *bent*, *soft*, *hard*, *narrow*, *broad*, and so on. The licensing of these metaphorical uses of adjectives can be represented via a

lexical rule, which is a more specific version of the Metaphor Lexical Rule, in that it instantiates the argument that the adjective modifies to be of type **human** (rather than $\top$). But this lexical rule would not specify the semantics of the metaphorical sense in any detail. Pragmatics will be needed to compute this.

Having captured this lexical generalisation, we can capture a further one: that physical objects can apply as predicates to humans, and they receive a metaphorical interpretation where the adjectives that apply to the physical object are now applied in their metaphorical sense to the human. So, for example, the lexical rule would predict that *rock* can take an argument of type **human**, and that in this metaphorical use, the adjectives that apply to the original **physical-object** meaning of *rock*, now apply in their metaphorically shifted sense to the **human**. So (13) is interpreted as *John is solid, heavy, hard to move*, and so forth.

(13) John is a rock.

It must be stressed, however, that since the metaphorical senses of the adjectives are underspecified in the lexicon, then so is the metaphorical sense of *rock*. The lexical rule licenses the predicate argument structure in (13), but gives few clues about the resulting meaning. Pragmatics reasoning is needed for this task. This is why although the interpretation of (13) is relatively clear (because this is a fairly well-used, and therefore according to Searle (1979), it is arguably an established use of *rock*), the metaphorical meaning of (14) is unclear, without further pragmatic information, even though the predicate-argument structure of the sentence is licensed by the grammar:

(14) ?Sam is a pebble.

We now turn to the task of modeling how pragmatics influences metaphorical interpretation, given the confines of metaphor specified by the Metaphor Lexical Rule (and its subtypes of lexical rules).

5 Metaphor in Context

The Metaphor Lexical Rule delimits the possibilities for metaphorical interpretation, but as is standard in constraint-based approaches to lexical

semantics, it leaves the meanings of words underspecified. On some occasions, e.g., in (5a) the semantic types of the arguments with which the words combine can serve to resolve some of this underspecification. But in general, pragmatic reasoning is required to flesh out the interpretation of metaphor in a discourse context. It is well-known that pragmatic information such as domain knowledge influences metaphorical interpretation (e.g., Lakoff and Johnson, 1980). We take this further, and claim that *rhetorical relations* such as *Narration* and *Contrast* also influence the meaning of metaphor.

Let's consider a particular example where discourse structure (as defined by rhetorical relations) and pragmatics influence the interpretation of metaphor. We observed in section 4 that the meaning of (14) in isolation of any discourse context is hard to compute. But it's acceptable in (15):

(14) ?Sam is a pebble.

(15) a. John is a rock.

b. But (compared to John) Sam is a pebble.

We think that the *Contrast* relation in (15) provides the information we need in order to calculate the metaphorical meaning of *pebble* at the discourse level. Note that it cannot simply be domain information about pebbles and rocks, since this was available for the interpretation of (14) and failed by itself to yield a coherent interpretation. Rather, the clue lies in the *juxtaposition* of (15a) and (15b), and the fact that (15b) *Contrasts* with (15a).

To see how the *Contrast* relation determines metaphor in (15), we'll link the lexical reasoning described above, to the semantics of *Contrast* in SDRT (Asher, 1993). First, we must build the logical forms of the two sentences compositional via the grammar. Because of the type clash in constructing the compositional semantics from the conventional lexical entries for (15a), we are forced to a metaphorical interpretation of *rock* using the more specific lexical metaphorical rule concerning **physical-objects**. For the sake of simplicity, we assume (15a) is simply represented by the logical form $\alpha(j)$, where α represents the semantic content of the metaphorical interpretation of the predicate *rock*.

We have argued that the metaphorical interpretation of *rock* is established, and therefore we assume here that α is computed to be *reliable*. However, the Metaphor Lexical Rule gives us further information. The value of

SEM:INDEX:QUALIA:FML, or in other words, the *differentiae*, is preserved in the lexical entry for the (metaphorical use of) *rock*.

Before we discuss the value of the differentiae for *rock*, let us produce the syntagmatic representation of (15b). After doing that, we must reason about how it attaches to the constituent $\alpha(j)$ with a rhetorical relation. As with (15a), the grammar fails to build the syntagmatic representation of (15b) out of the conventional lexical entries because of type clashes in the TFSs. We must use a metaphorical interpretation of *pebble*. In contrast to the semantic content of the first constituent $\alpha(j)$, however, we fail this time to obtain enough information in this process to get an informative constituent $\beta'(s)$.

Using the techniques described in Lascarides and Asher (1991, 1993) and elsewhere, we must attach the constituent $\beta'(s)$ through attaching this constituent with a rhetorical relation to the discourse context. In general, computing the rhetorical relation involves nonmonotonic reasoning on the reader's background knowledge. But in this case, the rhetorical relation must be *Contrast*, as indicated by the cue word *but*. So we must check the coherence of the *Contrast* relation in (15), by ensuring that the discourse constituents—which are the DRSs that represent the two sentences being connected together—are isomorphic both structurally and semantically, and that there is a contrasting theme between them. There is a default heuristic that one aims for as much structural and semantic isomorphism as possible. Maximum isomorphism is achieved, if *rock* is associated with *pebble* in the structural mapping between the constituents, and the semantics of these items contrast with each other.

In SDRT computing a rhetorical relation between constituents can trigger modifications to the semantics of the constituents themselves (Asher 1993, Asher and Lascarides 1994, 1996). Intuitively, this process corresponds to accommodating the information that's necessary to ensure discourse coherence is maintained (cf. Stalnaker 1968). It captures the intuition that speakers expect hearers to infer information that's not explicitly said during discourse processing. Each rhetorical relation imposes coherence constraints in terms of the semantic relation between the constituents being attached. If the necessary information isn't already there then it can be added in a constrained manner. Therefore, since we successfully interpret the metaphorical use of *rock* in (15a), we can use the coherence constraints on *Contrast* to infer further semantic content for the interpretation of *pebble* in this context.

Pebble will mean whatever produces the maximally contrasting theme, given the conventional constraints defined by the Metaphor Lexical Rule.

Now, we assume the conventional sense of *rock* has a value for SEM:QUALIA:FML which distinguishes it from other objects in the same semantic class (e.g., the conventional senses of *boulder*, *pebble* and so on). So, the conventional sense of *rock* specifies in SEM:QUALIA:FML that rocks are bigger than pebbles, but smaller than boulders. Similarly, the value of SEM:QUALIA::FML in the (conventional sense of) *pebble* contains the information that pebbles are smaller than rocks. So, given that the Metaphor Lexical Rule ensures that the values of SEM:QUALIA:FML (i.e., the differentiae) survive in metaphorical interpretation, these differentiae provide us with a scale with which we can compare (the metaphorical senses of) *pebble* and *rock*. The partial order of sizes between pebbles and rocks etc given in the differentiae becomes a scale of reliability, since we know *rock* means *reliable* in (15a). And since *pebble* precedes *rock* on this partial order, *pebble* means *less reliable*. Moreover, to maximise the contrast as dictated by the rhetorical relation *Contrast* between (15a) and (15b), *pebble* doesn't just mean *less reliable*; it means *unreliable* in this discourse context. On the other hand, *boulder* in (15a,c) means *more reliable*.

(15) c. But (compared to John) Sam is a boulder.

d. ?But (compared to John) Sam is a stone.

And since *stone* isn't specified according to size in the differentiae of the conventional entries, it doesn't feature in the partial order, and so one cannot compute the degree of reliability that *stone* is supposed to mean in (15a,d). In this case, the *Contrast* relation doesn't help us further specify the metaphorical meaning of *stone*. And so our theory correctly predicts that (15a,d) is anomalous.

SDRT's strategy of accommodation comes into play in cases like these. The need for discourse coherence triggered a modification of the truth conditional content of the representation $\beta'(s)$ of the meaning of (15b) (again, for details see Asher 1993), so that we obtain the necessary contrasting theme. Indeed, according to this revision technique as it's defined in Asher (1993), we replace $\beta'(s)$ with $\neg\alpha(s) \wedge \beta'(s)$.[6] This gives us the right intuitive results: namely,

[6] Note that given that β' is so underspecified, this conjunction is satisfiable. If it weren't satisfiable, then the *Contrast* relation would be incoherent.

that (15b) in this context means that Sam is a pebble, and not a rock. A similar inference holds for (15a,c).

We now consider an example where the discourse context actually triggers a metaphorical interpretation. Sentence (16), in the absence of information to the contrary, means that I (conventionally) climbed a (**physical-space**) greasy pole.

(16) I have climbed to the top of that greasy pole.

There is no reason to believe that the sentence is metaphorical, because the grammar permits the conventional sense of the CoL verb *climb* and the **physical-space** *greasy pole* to combine successfully. However, in (17)—which is a slightly expanded example from Searle (1979), that he attributes to Disraeli—the *Contrast* relation again helps us to determine the metaphorical interpretation of (17b). But further and in contrast with (15), the relation, together with the proposition expressed by the first sentence, *trigger* the metaphor in the first place.

(17) a. I have always despised politics.

b. But I have climbed to the top of that greasy pole.

Again using the techniques in SDRT, one must check the coherence of the *Contrast* relation, and aim for maximum structural and semantic isomorphism between the constituents. Furthermore, the anaphoric expression *that greasy pole* must be resolved to an available antecedent from the discourse context. Given the discourse context, SDRT predicts there is only one candidate antecedent: *politics*. But the type hierarchy in the constraint-based grammar prevents *that greasy pole* from being identified with *politics*, unless it undergoes type coercion. Therefore, using a principle of Charity familiar from philosophical work on metaphor (Black 1962, 1979, Searle 1979), we learn that *that greasy pole* is to be interpreted metaphorically, and is identified with politics (and so *greasy* is assigned a metaphorical interpretation, which is ascribed to politics). Thus, through calculating the anaphora resolution, with the aid of the discourse structure, we have triggered a metaphorical interpretation of *that greasy pole*.

But now, the conventional CoL verb *climb* cannot combine with *that greasy pole* anymore, because the CoL verb *climb* requires an argument of type

physical-space, and the TFS we must use to represent *that greasy pole*, which in turn has been chosen on the basis of constraints on anaphoric links in the discourse, is not of this type. Therefore, using the principle of Charity again, *climb* must be interpreted in a different way.

The constraints described in the Metaphor Lexical Rule, ensure respectively that the differentiae component is definitely preserved, and furthermore, that as many properties as possible, in addition to the *differentiae* of conventional *climb*, are preserved in the metaphorical interpretation. The former property means that all the properties of the source, intermediate path and goal of *climb* in (17b)—save that they are of type **physical-space**—are preserved. Similarly for the phrase *to the top*. So, Disraeli starts at the bottom of politics, and climbs to the top. We must therefore assign *politics* a qualitative scale: the career hierarchy (clerk to prime minister) is a plausible candidate, and under this interpretation, (17b) means that Disraeli worked his way up the career hierarchy. Note that the way we view the orientation of the career hierarchy is itself an example of metaphor, which is determined in part by the fundamental values in our culture (Lakoff and Johnson, 1980). Thinking of clerk as at the bottom and prime minister as at the top is a coherent metaphorical interpretation of the concept of career, whereas the opposite orientation, with the clerk at the top and prime minister at the bottom, would be incoherent.

A particularly compelling metaphor uses a combination of world knowledge and discourse structure. Arguably, this occurs in Romeo's metaphor for Juliet:

(18) a. What light through yonder window breaks?

b. It is the East, and Juliet is the sun!

William Shakespeare, Romeo and Juliet, Act II, Scene 2

Since (18a) is a question, by the constraints in SDRT it must form part of a question answer pair (*QAP*) (Asher and Lascarides, 1996). An answer to this question is a proposition which asserts what the light is. Indeed, as we'll see, (18b) gives us an answer, but not a direct one; we have to infer what the direct answer is.

One task is to resolve the anaphor *it* in (18b). Just as in (17), the discourse context constrains this to be identified with either the window or the light.

Both invoke type violations (for neither light nor windows can be the East, although they can be *in* the East), and will force a metaphorical interpretation. The identification of the anaphor with *window* is more plausible, because the mass term *light* doesn't determine a physical location to the extent that *window* does. So we have determined that the physical window is the metaphorical East. The second clause of (18b) informs us that Juliet is the metaphorical sun. This is computed along the same lines as the analysis of sentence (13), since we are associating a human with a natural object. Therefore, Juliet is ascribed with the adjectives, metaphorically interpreted, that conventionally apply to the sun, such as light and radiance. Typically, these are connected with beauty, and so we infer from the second clause of (18b) that Juliet is beautiful. However, we still have to determine how (18b) answers the question in (18a). Using the world knowledge that the sun rises in the East, one can assume that the metaphorical interpretation of (18b) is also one where the metaphorical sun rises in the metaphorical East. Since the metaphorical East is, physically, the window, and the metaphorical sun is Juliet, this information places Juliet at the window. Having learned that Juliet, who is the sun, is at the window, the question in (18a) is answered: Juliet is the light breaking at the window. One could infer further content, using world knowledge of the sort Shakespeare might have assumed his audience would have at their disposal (like a roughly Ptolomaic view of the universe); for example, as Juliet is the metaphorical sun, she is at the center of Romeo's world.

6 Conclusion

We have shown how to compute metaphorical interpretations from a combination of two mechanisms: lexical rules which specify the range of possible meaning shifts of classes of words; and a theory of discourse structure, which provides mechanisms for adding truth conditional content to metaphorically interpreted constituents, relative to the context in which they're uttered. We showed how this can explain data concerning: verbs involving change of location; the metaphorical shift of meaning of words that refer to kinds of physical objects when they are predicated of persons; and the dependence of metaphorical interpretation upon discourse structure. What we have done, of course, is very far from a comprehensive theory of metaphor. At best we have offered a proof of concept of an approach. But by using the modern

logical tools of formal pragmatics and semantics, we hope that we and others can make progress on this difficult subject and that in turn a better understanding of metaphor will enhance our understanding of lexical meaning and lexical processes.

7 Acknowledgements

This paper was presented at the 1995 AAAI Spring Symposium Workshop entitled *Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity and Generativity.* We would like to thank the participants of this workshop for their useful comments. In addition, we would like to thank Julius Moravscik for his comments on an earlier draft of this paper.

8 References

Asher, N. [1993] *Reference to Abstract Objects in Discourse*, Kluwer Academic Publishers.

Asher, N. and A. Lascarides, [1994] Intentions and Information in Discourse, *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*, pp34–41, Las Cruces.

Asher, N. and A. Lascarides, [1995] *Lexical Disambiguation in a Discourse Context*, *Journal of Semantics*, **12**, pp69–108.

Asher, N. and A. Lascarides, [1995b] Questions in Dialogue, HCRC research report number HCRC-81, available from HCRC, University of Edinburgh.

Asher, N. and P. Sablayrolles, [1995] A Typology and Discourse Semantics for Motion Verbs and Spatial PPs in French, *Journal of Semantics*, **12**.

Aurnague, M. and L. Vieu, [1993] Towards a Formal Representation of Space in Language: A Commonsense Reasoning Approach, *Proceedings of the* IJCAI93 *workshop on Spatial Reasoning.*

Black, M. [1962] *Models and Metaphors*, Cornell University Press.

Black, M. [1979] More About Metaphor, in Ortony, A. (ed.) *Metaphor and Thought*, Cambridge University Press, 1979, pp19–45.

Briscoe, E. J. and A. Copestake [1997] Lexical Rules in the TDFS Framework: The Dative Construction, submitted to *Computational Linguistics*.

Briscoe, E. J., A. Copestake, and B. Boguraev, [1990] Enjoy the Paper: lexical semantics via lexicology, *Proceedings of the 13th International Conference on Computational Linguistics (COLING90)*, Helsinki, pp42–47.

Boguraev, B., and J. Pustejovsky, [1990] Lexical Ambiguity and the Role of Knowledge Representation in Lexicon Design, *Proceedings of the 13th International Conference on Computational Linguistics (COLING90)*, pp36–42, Helsinki.

Copestake, A. and E. J. Briscoe [1992] Lexical Operations in a Unification Based Framework, in J. Pustejovsky and S. Bergler (eds.) Lexical Semantics and Knowledge Representation. Proceedings of the first SIGLEX Workshop, Berkeley, CA, pp101–119, Springer-Verlag, Berlin.

Copestake, A. and E. J. Briscoe [1995] Semi-Productive Polysemy and Sense Extension, *Journal of Semantics*, **12.1**, pp15–67.

Davidson, D. [1984] What Metaphors Mean, in D. Davidson *Inquiries into Truth and Interpretation*, pp245–264, Clarendon Press.

Evans, R. and G. Gazdar, [1989a] Inference in DATR, *Proceedings of the 4th Conference of the European Chapter of the Association for Computational Linguistics (EACL89)*, Manchester, England, pp66–71.

Evans, R. and G. Gazdar, [1989b], The semantics of DATR, in Cohn, A. G. (ed.) *Proceedings of the Seventh Conference of the Society for the Study of Artificial Intelligence and Simulation of Behavior (AISB)* pp79–87, Pitman/Morgan Kaufmann, London.

Flickinger, D. [1985] *Lexical Rules in the Hierarchical Lexicon* PhD thesis, Stanford University.

Hesse, M. [1966] *Models and Analogies in Science*, University of Notre Dame Press.

Lascarides, A. and N. Asher, [1993] Temporal Interpretation, Discourse Relations, and Commonsense Entailment, *Linguistics and Philosophy* **16**, pp437–493.

Lascarides, A., E. J. Briscoe, N. Asher, and A. Copestake [1996] Order Independent and Persistent Default Unification, *Linguistics and Philsophy*,

19(1), pp1–89.

Lascarides, A. [1995] The Pragmatics of Word Meaning, *Proceedings of the* AAAI *Spring Symposium Series: Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity and Generativity*, pp75–80,

Lascarides, A. and A. Copestake, [1997a] The Pragmatics of Word Meaning, to appear in *Journal of Linguistics*.

Lascarides, A. and A. Copestake [1997b] Default Representation in Constraint-based Frameworks, submitted to *Computational Linguistics*.

Lascarides, A., A. Copestake, and E. J. Briscoe [1996] Ambiguity and Coherence, *Journal of Semantics* **13**(1), pp41–65.

Lakoff, G. and M. Johnson [1980] *Metaphors We live By*, The University of Chicago Press.

Mereurs, D. [1995] Towards a Semantics for Lexical Rules as Used in HPSG, *Proceedings of the* ACQUILEX *Workshop on Lexical Rules*, Cambridge.

Moravscik, J. M. [1990] *Thought and Language*, Routledge, London.

Moravscik, J. M. [this volume] Metaphor, Creative Understanding and the Productive Lexicon.

Ortony, A. [1979] *Metaphor and Thought*, Cambridge University Press.

Pollard, C. and Sag, I. A. [1994] *Head-Driven Phrase Structure Grammar*, CSLI, University of Chicago Press, Chicago and London.

Pustejovsky, J. [1991] The Generative Lexicon, *Computational Linguistics* **17** 4, pp409–441.

Pustejovsky, J. [1995] *The Generative Lexicon*, MIT Press, Cambridge Mass.

Searle, J. [1979] Metaphor, in Ortony, A. (ed) *Metaphor and Thought*, Cambridge University Press, pp92–123.