# Context sensitive reasoning with lexical and world knowledge

**Anthony Hunter, Department of Computer Science, UCL &**
**Lutz Marten, Department of Linguistics, SOAS**

## 1. Introduction

The question of polysemy and the representation of different lexical word-senses has recently attracted an increased interest in the linguistics and computational linguistics literature. In particular work in Generative Lexicon Theory (GLT, cf. Pustejovsky 1995, 1998, Lascarides, Copestake & Briscoe 1996, Lascarides & Copestake 1998) has proposed to enrich the format of lexical representation so that a number of 'novel' word senses can be generated. In this paper, we discuss GLT and show that, while being able to derive novel word-senses, there are, however, problems with this approach. In particular, we argue that GLT lacks the expressivity necessary for supporting context sensitive reasoning, and show this by considering relevant examples involving underspecification inherent in verbs and the process of establishing semantic predicate-argument structures in utterance interpretation. The theoretical background of our work is the notion of utterance interpretation as explored in Relevance Theory (Sperber & Wilson 1986/95) and Dynamic Syntax (Kempson, Meyer-Viol, and Gabbay 1999), and in particular the idea of verbal underspecification developed in Marten (1999), where it is argued that verbs are lexically underspecified as to their semantic type, so that a verb's eventual arity will only be established in the utterance context. From this perspective, the interplay between this structural underspecification, lexical polysemy and context sensitive interpretation leads to an alternative view on knowledge representation employed in utterance interpretation and parsing, namely that hearers employ world knowledge, rather than 'lexical' knowledge, in the task to derive the proposition expressed, and that enrichment of meaning involves a level of conceptual representational structure. In this paper, we are not discussing the theoretical aspects of this view in detail (cf. Hunter & Marten 1999, Marten 1999, and also Sperber & Wilson 1999?, Fodor & Lepore 1998?, Carston 1996) but rather we propose a formalization of enrichment of natural language expressions as used in parsing by means of default logic statements (Reiter 1980). We set out by discussing GLT in some detail, and then sketch our formalization of the basic architecture for relating parsed natural language structures and world knowledge.

## 2. The Generative Lexicon

Generative Lexicon Theory (GLT) is one recent formulation of lexical knowledge. It is part of a wider trend in (computational) linguistics to analyse lexical knowledge not as an ideally minimal set of knowledge specifying only basic syntactic and maybe thematic information, but rather to think of lexical items as more structured, for example as being related by semantic relations modelled as feature structures, or as specifying certain aspectual information (cf. e.g. Sag & Szabolcsi 1992, Pustejovsky & Boguraev 1996). In this section, we introduce GLT as formulated in Pustejovsky (1995) and discuss an extension of the theory proposed in Lascarides, Copestake & Briscoe (1996) and Lascarides & Copestake (1998).

### 2.1. Generative Lexicon Theory

GLT as originally formulated by Pustejovsky (1995) proposes that information provided by lexical items is much richer than standardly assumed. In particular, GLT argues that the lexicon of a given natural language cannot simply be characterized as a list of items with only syntactic and minimal semantic information, for three reasons (1995: 39):

1. Words can be used creatively; they assume new senses in novel contexts.
2. Word senses are not atomic; they overlap and refer to other senses of the word.
3. A single word sense can have multiple syntactic realizations.

The creative use of words is found, for example, with the adjective *fast* (1995: 44/45):

(1a) a fast boat
(1b) a fast typist
(1c) a fast book
(1d) a fast driver
(1e) a fast decision
(1f) a fast motorway

The meaning of *fast* varies according to the noun it modifies, thus, Pustejovsky argues, giving rise to at least four different lexical entries (1995: 44/45):

(2a) **fast$_1$**: to move quickly
(2b) **fast$_2$**: to perform some act quickly
(2c) **fast$_3$**: to do something that takes little time
(2d) **fast$_4$**: to allow for driving a vehicle quickly on it
(ie. the motorway)

There is seemingly no principled end to such a list, nor any way to relate one sense to another. Pustejovsky (1995: 45/46) presents the examples in (3) as being indicative of a similar problem; since a lexical entry for *want* would have to include at least the ones found in (4):

(3a) Mary wants another cigarette.
(3b) Mary wants a beer.
(3c) Mary wants a job.

(4a) **want$_1$**: to want to smoke
(4b) **want$_2$**: to want to drink
(4c) **want$_3$**: to want to have

Again, an increase in NP objects of *want* would lead to an increase in lexical entries (in addition to entries required due to varying syntactic complementation).

The problem of overlapping of different word senses can be seen with *bake*:

(5a) John baked the potatoes.
(5b) John baked a cake.

The verb *bake*, along with *cook* and *fry*, is, according to GLT, ambiguous between a 'change-of-state' and a 'creation' reading; in (5a), the activity of baking changes the state of the object, potatoes, from cold to hot (and not_really_edible to edible), while in (5b), the object is the result of the activity – it did not previously exist. However, according to GLT, these two senses are not clearly distinguishable, even in context, since one sense is included in the other, and since the actual activity is not distinguished by the different objects. Thus, even if one were to allow many different senses, the postulation of two senses for *bake* here would obliviate their partial overlap.

Finally, Pustejovsky observes that a single form can participate in a number of syntactic realizations, corresponding to different senses (1995: 51):

(6a) Madison Avenue is apt to forget that most folks aren't members of the leisure class. (factive reading)

(6b) But like many others who have made the same choice, he forgot to factor one thing into his plans: Caliphobia.
(non-factive reading)

(6c) As for California being a state being run by liberal environmental loonies, let's not forget where Ronald Reagan came from.
(embedded question)

(6d) What about friends who forget the password or never got it?
(concealed question)

(6e) He leaves, forgets his umbrella, and comes back to get it.
(ellipsed non-factive)

By simply postulating several entries for *forget*, both the relation between complement type and reading, as well as the common 'core meaning' of the verb are not expressed.

Thus, Pustejovsky concludes that the only way for a list-lexicon to deal with these phenomena is the postulation of an infinity of different senses for a single lexical item, which is not only unintuitive and cumbersome, but also fails to capture any structure or systematicity between several senses, both semantic and syntactic.

In contrast, GLT designs lexical entries which are structured, and thus can encode different senses, and which allow for the generation of novel word senses by composition of different features. That is, the different readings of *bake* (as in (5)) result from the complex lexical entry of the verb and from information encoded in the lexical entry of the object. It is this process of generativity, devising new senses by composition, by which GLT proposes to meet the demands of accounting for creativity, overlap of senses, and multiple syntactic realizations.

In order to achieve this, GLT introduces four levels of lexical semantic representation (1995: 61):

- argument structure
- event structure
- qualia structure
- lexical inheritance structure

Argument structure states in a standard way number and type of logical arguments, with the addition of 'shadow' and 'default'-arguments, different types of semantically necessary, but syntactically optional arguments. Event structure encodes aspectual lexical information, similar to, but more refined than, a simple event-variable. Qualia structure is probably the most novel (and controversial) idea in GLT, and it is also the structure with the most complicated internal structure. Qualia structure values encode information about what the lexical item is (refers to), what it consists of, what it is made of, and what it is used for. The relation among lexical items is encoded in the final structure, lexical inheritance, which expresses hyponymy and other lexical relations on a particular lattice structure. Information from individual lexical items is related to the lexical inheritance structure via the qualia and formal structure values of the lexical item. Similar to HPSG, values of predicates can be co-indexed to indicate feature unification. This particular selection of features is partly motivated by the claim that the information included in the entries is relevant for speakers' knowledge of language; it is claimed to play a role in grammar which distinguishes it from (other) world knowledge[1].

The lexical entry for *bake* thus looks as in (7) (1995: 123):

(7)

$$\begin{bmatrix} \textbf{bake} & \\ \text{EVENTSTR} = & \begin{bmatrix} \text{E}_1 = \textbf{e}_1\textbf{:process} \\ \text{HEAD} = \textbf{e}_1 \end{bmatrix} \\ \text{ARGSTR} = & \begin{bmatrix} \text{ARG1} = \boxed{1}\begin{bmatrix}\textbf{animate\_ind} \\ \text{FORMAL} = \textbf{physobj}\end{bmatrix} \\ \text{ARG2} = \boxed{2}\begin{bmatrix}\textbf{mass} \\ \text{FORMAL} = \textbf{physobj}\end{bmatrix} \end{bmatrix} \\ \text{QUALIA} = & \begin{bmatrix}\textbf{state\_change\_lcp} \\ \text{AGENTIVE} = \textbf{bake\_act}(\,\textbf{e}_1, \boxed{1}, \boxed{2}\,)\end{bmatrix} \end{bmatrix}$$

The entry illustrates possible values to the three structures of lexical information encoded in lexical entries (the fourth structure, lexical inheritance, serves as ontological backbone which helps interpret the values of the FORMAL parameter). The event structure value identifies 'baking' as a process, the headedness value is exploited for complex event structures and is here of minor importance. Argument structure makes the verb transitive and places further restrictions on the arguments[2]. The qualia structure values, finally, identify *bake* as belonging to the 'change of state' 'lexical-

---

[1] Which seems to imply that, unless thinking is construed as being dependent on language, some information is being duplicated since for example the observation that one can eat cakes is part of the lexical meaning of *cake*, but surely we know that independent of the word *cake*, so that both the lexical item and the world knowledge include this statement.

[2] There is actually some mismatch here between what is given in the typed feature structure and what, in our reading of the surrounding text, should have been in there; ARG2 should probably be a default or shadow argument to include sentences like 'John was baking in the kitchen'; furthermore, it is unclear why ARG2 is 'mass' here, since, as discussed below, the default reading of *bake* for Pustejovsky is the baking of things like potatoes, hence it should be 'count' (cf. also Fodor & Lepore (1998)).

conceptual paradigm' (lcp), over which generalizations over verb-classes can be stated. 'Agentive' means that the (saturated) predicate involves some 'bringing about'. In GLT, this is the only lexical entry for *bake*. The claim is now that objects such as *potatoes* or *carrots* leave the information from *bake* unmodified, but that objects like *cake* 'shift' the reading of *bake* to a resultative or 'create' reading. That is, apparently different senses of the verb really arise from interaction with the lexical specification of (object) NPs, in particular by co-composition of qualia values. The process is triggered from object NPs involving words like *cake* (1995: 123):

(8)

$$\begin{bmatrix} \textbf{cake} & \\ & \\ \textsc{argstr} = & \begin{bmatrix} \textsc{arg}1 = \textbf{x:food\_ind} \\ \textsc{d-arg}1 = \textbf{y:mass} \end{bmatrix} \\ & \\ \textsc{qualia} = & \begin{bmatrix} \textsc{const} = \textbf{y} \\ \textsc{formal} = \textbf{x} \\ \textsc{telic} = \textbf{eat}(\textbf{e}_2, \textbf{z}, \textbf{x}) \\ \textsc{agentive} = \textbf{bake\_act}(\textbf{e}_1, \textbf{w}, \textbf{y}) \end{bmatrix} \end{bmatrix}$$

The 'D-ARG' (i.e. default argument) indicates that cakes are made from stuff which can optionally be expressed as an oblique argument as for example in *bake a cake from/with flour*. This is taken up in the qualia structure, where CONST states that cakes consist of (are constituted by) the stuff encodable as a default argument, dough, maybe, or flour, or chocolate (but, not apparently, eggs, which are count). FORMAL encodes most directly what the word actually means, namely food (by feature sharing with ARG1), and this value can be found again at the lexical inheritance structure to give the embedding in the lexical net. TELIC means function, here that one eats cakes, while AGENTIVE indicates how the cake comes into existence, namely by baking. Note that the AGENTIVE value has the same predicate as *bake*, and that the distribution of the variables is such that the eating event involves an as yet unknown eater (z) and the food variable from ARG1 (x), but that baking involves not (x) in object position, but (y), the variable bound as 'mass' in D-ARG1 – that is, the act of baking does not involve the cake, but rather the stuff out of which it is made. Now the combination of *bake* with *a cake* results in the following semantic representation (1995: 125):

(9)

$$\left[\begin{array}{ll}
\textbf{bake a cake} & \\
 & \\
\textsc{eventstr} = & \left[\begin{array}{l} \text{E}_1 = \mathbf{e_1 : process} \\ \text{E}_2 = \mathbf{e_2 : state} \\ \text{R}\textsc{estr} = <_{\alpha} \\ \text{H}\textsc{ead} = \mathbf{e_1} \end{array}\right] \\
 & \\
\textsc{argstr} = & \left[\begin{array}{ll} \textsc{arg}1 = \boxed{1} & \left[\begin{array}{l} \textbf{animate\_ind} \\ \textsc{formal} = \textbf{physobj} \end{array}\right] \\ & \\ \textsc{arg}2 = \boxed{2} & \left[\begin{array}{l} \textbf{artifact} \\ \textsc{const} = \boxed{3} \\ \textsc{formal} = \textbf{physobj} \end{array}\right] \\ & \\ \textsc{d-arg}1 = \boxed{3} & \left[\begin{array}{l} \textbf{material} \\ \textsc{formal} = \textbf{mass} \end{array}\right] \end{array}\right] \\
 & \\
\textsc{qualia} = & \left[\begin{array}{l} \textbf{create-lcp} \\ \textsc{formal} = \mathbf{exist(e_2,} \boxed{2}\ \mathbf{)} \\ \textsc{agentive} = \mathbf{bake\_act(e_1,} \boxed{1}, \boxed{3}\ \mathbf{)} \end{array}\right]
\end{array}\right]$$

Amongst other details, the object results in a modified qualia structure which makes the VP (not the verb, incidentally) a member of the 'create' lcp. Furthermore it states that there are two events - one, the AGENTIVE value, involving the baker and the cake ingredients, and a second one, the FORMAL value, which is a stative event (cf. the event structure value for $\text{E}_2$) at which the second argument, the cake, exists. Of interest are that several semantic intuitions are captured at the interrelated but distinct lexical levels; thus the resulting state of existence can be read off from the event and the qualia structures, while the interplay between argument and qualia structures allows for a mismatch between syntactic and semantic structure – *a cake* is the syntactic direct object, but is not a member of the extension of **bake_act** in the QUALIA.

GLT has highly structured lexical information for members of all parts of speech, and claims with those to provide a principled explanation for different (for GLT, lexical) meanings of words in context, accounting for a large range of readings, as well as for why certain readings are impossible, without the necessity for multiple lexical entries.

## 2.2. The Problem of Context Sensitivity

It is clear from the exposition of GLT so far that this approach assumes that there is a distinct and rich grammatical level of lexical information. A number of features and feature values are part of the encoded meaning of lexical items. However, under this conception one problem immediately arises. This is the demarcation of lexical and world knowledge, and the attendant problem of the demarcation of (the relevant) context (cf. Sperber & Wilson 1997). Since the GLT entries, though rich, are restricted to information which is 'grammatically relevant', only some contextual differences in word meaning are expressed. Thus for example, since the difference between the change of state and the create senses of *bake* is partly tied to the mass/count distinction of the (D-argument of the) object, there is no difference between (10) and (11):

(10) John was baking a cake.
(11) John was baking a bread.

While the concepts involved here are probably similar, there still is the possibility to construe them differently on occasion; for example, the baking of bread might involve more preparation and time, different ingredients, different tools, etc. than the baking of cakes[3]. Yet this difference is not expressed in GLT, which in fact predicts that the two senses of *bake* are completely identical.

On the other hand, since the creative reading results from the unification of qualia values, in particular only from those where the object's qualia value explicitly states that the object is created by baking, all other cases fall under the change of state reading, which is provided in the verb's entry. Thus, the meaning of (12) and (13) come out as the same:

(12) John was baking a potato.
(13) John was baking a flower.

But, as in the examples above, there are contexts where the baking of potatoes might differ from the baking of flowers. Furthermore, the change of state reading should be available for objects which are made by baking:

(14) John was baking a pizza.

In a context where the pizza is deep-frozen, *bake* has no creative meaning, but given the way qualia structures unify, this reading is not obtainable, unless an ambiguity between pizzas and deep-frozen pizzas is postulated, which is what GLT is trying to avoid.

All the examples discussed so far point to problems with GLT which the theory according to its own aims should be able to handle. More general problems of this approach, as for example discussed in Fodor & Lepore (1998), include that GLT specifies just an arbitrary subset of world knowledge we have about things in the world, rather than about words. In principle, according to this criticism, there is nothing to be gained from writing the fact that cakes are made by baking into the lexicon. Another problem concerns the use of typed feature structures for the representation of this kind of knowledge; feature unification as used in GLT does not support any logical reasoning, all lexical items have to be fully specified and be assigned a place in the feature hierarchy. Furthermore, context sensitivity has to be specified in advance as the presence or absence of particular features. In this sense, the system is too restricted for modelling context sensitive reasoning with world knowledge.

The lack of context sensitivity and inferencing in GLT has been addressed by Lascarides, Copestake & Briscoe (1996) and Lascarides & Copestake (1998). They provide a conditional logic which models commonsense entailment and which

---

[3] A nice way to make bread is to use sour dough baked on low heat in a wood fired stone oven. This doesn't work for cakes.

interacts with the senses provided by GLT[4]. Under this view, the novel senses of GLT, derived by feature unification, are default predicates which are part of the meaning of the sentence in which they occur. The second step of interpretation takes these typed feature structures as input and combines them into a model of discourse processing in which notions of discourse coherence (e.g. 'elaboration', 'contrast') are defined. The combination of GLT with discourse processing is exploited to provide an interpretation of default senses. A default sense can be checked against the world knowledge base. For the pizza example in (14) above, for example, the default information that John created the pizza results from the lexical information that pizzas are created by baking. The world knowledge base then might specify that baking deep frozen pizzas implies merely a state of change. This information is more specific (pertaining to more specific pizzas) than the information about pizzas in general. Now the conditional logic specifies that more specific information overrides more general information, and thus the default interpretation does not survive under embedding into the discourse, and *bake* is (again) interpreted in its change of state sense. In the absence of contradicting or more specific information, lexical defaults are taken over into the discourse. The model thus provides a means to combine the lexical senses from GLT with world knowledge and furthermore offers some indication of how the two interact.

There are, however, still problems with this enriched version of GLT. First, the model inherits the problems of GLT noted above, namely that the different senses are not fine-grained enough to distinguish different occasion specific concepts; different concepts with identical features (such as examples (10) - (13) above) are not distinguished at the discourse level, at least not if the defaults survive[5]. Secondly, this conception implies a rather unintuitive division of labour between lexical features and default reasoning; in the pizza example, the information that pizza baking (usually) means creating the pizza is lexical, while the information that baking a deep frozen pizza (usually) means changing the state of the pizza is part of the world knowledge, but the kind of information conveyed by the two statements is intuitively rather similar.

In conclusion, it appears that GLT, even when adjusted and amended, is not very suitable to model the process of concept formation relevant for the interpretation of underspecified predicate-argument structures. The approach suffers, from the perspective adopted here, from a suboptimal division of the notions of lexical and world knowledge which results in counter-intuitive analyses[6]. In the next section, we discuss how these problems can be overcome.

## 3. Default reasoning with world knowledge

In this section we introduce an alternative approach to the knowledge hearers employ in parsing, namely a default logic characterization of world knowledge, and show how the logic can be employed in the establishment of predicate argument structures. We assume here that reasoning applies to the output of the parser, that is, an account is

---

[4] This work is discussed more extensively in Hunter & Marten (1999) and Wilson (1999).

[5] If the default interpretations are overridden, more information might be included, depending on the world knowledge base. In order to model enrichment in the sense assumed in this paper, defaults would always be overridden and thus be pointless.

[6] There are also a number of potential problems associated with the use of conditional logic with defaults for this kind of reasoning, which might be better formalized by employing default logic; cf. Hunter & Marten (1999: 30/31) for further discussion.

given of how world knowledge interacts with predicates of varying arity, but not of how these predicates are formed. However, the dynamics of this process could possibly be added into the picture at a later stage. Furthermore, there is no prioritization of assumptions, that is, the notion of relevance remains unanalysed in the version discussed here. The work thus is mainly concerned with modelling reasoning with world knowledge and predicate argument structures under the assumption that lexical items address concepts as discussed in the following subsection. Consequently, all reasoning with linguistic structures is located in the world knowledge, which means in particular that it is open to logical inferencing and does not run over typed-feature structures. As a further contrast to GLT, we represent world knowledge as statements in default logic (as opposed to conditional logic) which incorporates classical first-order logic and provides an expressive and clear format for stating context sensitive reasoning[7]. We first introduce the system and compare it with the analysis given in section 2.

### 3.1. A default logic approach

In our system, the output of the parser is represented as function symbols which are translated into the logical language with recourse to world knowledge so that information from the parser and contextual default reasoning interact. Semantic information is represented as 'concepts', where a concept is a term of the logical language and accesses a set of logical statements (i.e. 'assumptions'), although the actual set accessed is at present left unspecified. Both these representations are part of the logical language (the world knowledge) expressed by default rules, so that the term accessing the set of assumptions may be part of assumptions in the set[8]. Reasoning with the output of the parser and world knowledge is characterized as comprising three related activities; commonsense checking, commonsense inferencing, and commonsense explaining.

Commonsense checking establishes whether a given output from the parser is consistent with world knowledge facts. The notions of normal and unusual information can be checked with reference to a (arbitrary) subset of default rules which might be called Normality default rules. For example, the idea that buttering toasts is usually done with a knife can be expressed with the following default rule:

(15) $$\frac{\text{Butter(X, the\_toast) : Butter(X, the toast, with(a\_knife))}}{\text{Butter(X, the\_toast, with(a\_knife)}}$$

The rule in (15) is a default rule. It states that, given Butter(X, the_toast) (i.e. the expression to the left of the colon, the 'precondition'), and furthermore given that Butter(X, the_toast, with(a_knife)) (i.e. the expression to the right of the colon, the 'jusitfication') is consistent with world knowledge, then Butter(X, the_toast, with(a_knife)) (that is now the expression under the line, the 'consequent') can be inferred[9]. The rule is stated over expressions of the logical language, not over the output of the parser directly. It is thus a piece of world knowledge, expressing an aspect of what (most) people know about buttering toasts. With reference to this rule,

---

[7] Cf. Reiter (1980) for a description of default logic.

[8] Cf. also Rips (1995) for a similar conception of psychological concepts.

[9] Technically, the inference is valid if no inconsistencies result with respect to the inferences derived, not with the total set of assumptions. It is presently left open which assumptions constitute relevant world knowledge.

the sentence in (16) can be checked, which we assume is represented as (17) after parsing:

(16) John buttered the toast with a knife.
(17) butter(John, the_toast, with(a_knife))

The first step in the interpretation of (17) is to translate the 'lexical' function symbol butter (lower case) in (17) into the world knowledge predicate Butter (upper case), which is formally achieved by the world knowledge predicate Holds applying to the output of the parse:

(18)
$$\frac{\text{Holds(butter(X, Y, Z)) : Butter(X, Y, Z)}}{\text{Butter(X, Y, Z)}}$$

The Translation default rule in (18) effectively states that the lexical predicate *butter* can be interpreted as the conceptual predicate Butter if it is consistent with the world knowledge to do so. Furthermore, commonsense inferences can be stated over the output from parsing using Holds. For example:

(19)
$$\frac{\begin{array}{c}\text{Holds(butter(X, Y, Z)) \& PartOf(X, X') \& PartOf(Y, Y') \&} \\ \text{Human(X') \& Food(Y') : Holds(butter(X, Y)}\end{array}}{\text{Holds(butter(X, Y))}}$$

The Facilitation rule in (19) allows the inference from the three place function symbol *butter* to the two place function symbol *butter* under the assumption that subject and object are human and food respectively, and that the inference is consistent with world knowledge. Holds then translates the inferred expression into a world predicate:

(20)
$$\frac{\text{Holds(butter(X, Y)) : Butter(X, Y)}}{\text{Butter(X, Y)}}$$

Now the inferred expression in (20) can be used as a precondition of the Normality rule in (15). The information in (17) is thus not only consistent with the world knowledge, but can also be seen as redundant in the sense that the same information can be inferred by Normality rules.

Before proceeding, we briefly point out what this system of reasoning does. This model assumes that all semantic information is located in the world knowledge, which is represented as a (large) set of first-order logic statements and default rules. The output of the parser is taken to be uninterpreted. Interpretation is provided by taking this output, that is here predicate–argument structures of varying arity, and relating it to the predicate symbols which are in the world knowledge. Given the complete generality of the system, both enrichment (as in (15)) and inferences (as in (19)) can be stated. Although the system requires a large number of individual statements and rules, it has the advantage that steps of interpretation are explicitly and clearly statable. The following discussion includes further examples of how world knowledge interacts with natural language representations.

The example considered above involved information consistent with world knowledge. The following example, in contrast, is unusual:

(21) butter(John, the_toast, with(a_spade))

By using Translation rule (18), the lexical predicate is translated into a world predicate:

(22) Butter(John, the_toast, with(a_spade))

Furthermore, by the rules in (19) and (20), the proposition in (23) can be inferred from (21):

(23) Butter(John, the_toast)

The assumption in (23) can in turn be used with the Normality default rule in (15) to give (24):

(24) Butter(John, the_toast, with(a_knife))

Under the assumption that the world knowledge includes the information that (22) and (24) are contradictory, the information in (21) can be flagged as unusual with respect to world knowledge. Note that the system merely indicates why (21) is unusual, but does not imply any resolution, which has to be stated separately. This contrasts with the GLT position where principles regulate the interpretation of defaults, e.g. in this example, the default inference (if it was arrived at) would be suppressed given that *with a spade* is part of the sentence. The aim here is more modest, since inconsistency is checked, but not resolved.

Commonsense explaining involves deriving new predicates, either world knowledge predicates or Holds inferences, so that commonsense inferencing as in (19) above can be seen as an instance of commonsense explaining. Another instance of commonsense explaining involves term substitution. By using term substitution, function symbols can be replaced by other, possibly more complex, function symbols, which might be more meaningful in a given context. That is to say, substitution licenses the translation of one predicate into another, or possibly several other predicates if doing so is consistent with world knowledge. For example, for the n-ary function symbol flies($\alpha_1, \ldots, \alpha_n$), where $\alpha_1, \ldots, \alpha_n$ are terms, the following substitutions might be useful:

(25a) Sub(flies($\alpha_1, \ldots, \alpha_n$), moves(through_the_air($\alpha_1, \ldots, \alpha_n$)))
(25b) Sub(flies($\alpha_1, \ldots, \alpha_n$), moves(through_a_trajectory($\alpha_1, \ldots, \alpha_n$)))
(25c) Sub(flies($\alpha_1, \ldots, \alpha_n$), moves(quickly($\alpha_1, \ldots, \alpha_n$)))
(25d) Sub(flies($\alpha_1, \ldots, \alpha_n$), moves(swiftly($\alpha_1, \ldots, \alpha_n$)))

Which one of possible substitution rules is appropriate is of course context dependent. Consider the examples in (26) and (27):

(26) flies(this_helicopter, to(the_island))
(27) flies(time, like(an_arrow))

These two examples can be rewritten using substitution with the following default rules (where T stands for zero or more further term variables):

(28) $$\frac{\text{PhysicalObject(X) : Sub(flies(X, T), moves(through\_the\_air(X, T)))}}{\text{Sub(flies(X, T), moves(through\_the\_air(X, T)))}}$$

(29) $$\frac{\neg\text{Aircraft(X) : Sub(flies(X, T), moves(quickly(X, T)))}}{\text{Sub(flies(X, T), moves(quickly(X, T)))}}$$

Application of these rules (and assuming the relevant world knowledge facts, e.g. PhysicalObject(this_helicopter)) results in the commonsense explanation of (26) and (27) as (30) and (31):

(30) moves(through_the_air(this_helicopter, to(the_island)))
(31) moves(quickly(time, like(an_arrow)))

As these examples show, term substitution thus provides another means of using default inferences in reasoning with parsed natural language strings and world knowledge. It should be noted that there is considerable overlap between substitution and Holds inferences; the information provided by (30) and (31), for example, could equally have been arrived at by a default rule using the Holds predicate. This adds to the expressivity of the system, which can be constrained according to application.

Whilst the system outlined in here does not provide a formal analysis of concept formation, there are a number of traits which make it an attractive starting point for the development of such an analysis. As the examples discussed above show, the system handles forms of enrichment. For example, the enrichment of a constituent, that is the inference from (32) to (33) is expressible:

(32) John buttered the toast.
(33) John buttered the toast with a knife.

Similarly, the inference from (34) to (35):

(34) John buttered the toast with a fork.
(35) John buttered the toast.

Furthermore, by substitution, (36) can be translated to (37):

(36) John buttered the toast.
(37) John applied an even layer of butter to the toast.

These three inferences are stated within one system of knowledge representation, which is capable of formalizing context sensitive and uncertain reasoning. By using explicit facilitation and translation rules, inferences can be incrementally added or retracted according to context. Furthermore, prototypical information can be represented by individually specified normality defaults. The expressive power thus exceeds typed feature structures and lattice theoretic representations, and provides a better means to represent processes of general reasoning.

To summarize, the system advocated here aims at providing a clear logic based formal account of how general reasoning interacts with output from parsing, in particular with predicate–argument structures of varying arity. The system is formulated in Default Logic, so that steps of inference are represented as default rules which license the commonsense checking, inferencing, and explaining of logical input structures, while being able to express context sensitivity and uncertainty. The underlying assumption is that natural language expressions address concepts, that is expressions of the logic, and that all interpretation is inferential, expressible in the system. The approach thus overcomes the problem of postulating two distinct levels of interpretation, such as seen with feature structures and conditional logic in GLT. The system is furthermore very expressive, since it offers several ways to derive particular information and can as such be suitably restricted given the need of particular applications.

## 4. Conclusion

In this paper we have discussed aspects of the formal representation of lexical and world knowledge. We have discussed Generative Lexicon Theory and extensions thereof, and have argued thats that approach is not fully satisfactory for the representation of context sensitive reasoning in utterance interpretation. We then have proposed an alternative formalization which does not assume rich lexical representations, but rather provides a direct interface between underspecified parsed strings and conceptual or world knowledge. We have shown how this interface can be modelled with recourse to default logic, and how a number of common processes at the interface, such as inferencing, enrichment, and inconsistency checking, can be modelled. Inevitably, the system is presented only in outline[10]. Yet we hope to have shown that our formalization provides a genuine alternative to typed-feature structure approaches, which has – once properly developed – both theoretical and empirical advantages.

## References

Carston, Robyn, 1996, 'Enrichment and loosening: complementary processes in deriving the proposition expressed', *UCL Working Papers in Linguistics* 8, 61-88.

Fodor, Jerry A. & Ernie Lepore, 1998, 'The emptiness of the lexicon: critical reflections on J. Pustejovsky's The Generative Lexicon', *Linguistic Inquiry* 25.

Hunter, Anthony & Lutz Marten, 1999, 'Reasoning with output from parsing using world knowledge', draft ms., UCL/SOAS.

Kempson, Ruth M., Wilfried Meyer-Viol & Dov Gabbay, 1999, *Dynamic Syntax*, (working title), ms., SOAS & King's College, University of London.

Lascarides, Alex & Ann Copestake, 1998, 'Pragmatics and word meaning', *Journal of Semantics* 11, 41-65.

Lascarides, Alex, Ann Copestake & Ted Briscoe, 1996, 'Ambiguity and coherence', *Journal of Linguistics* 34, 387-414.

Marten, Lutz, 1999, *Syntactic and Semantic Underspecification in the Verb Phrase*, doctoral disseratation, SOAS, Lonodn.

Pustejovsky, James & Branimir Boguraev, eds., 1996, *Lexical Semantics*, Oxford: Clarendon.

Pustejovsky, James, 1995, *The Generative Lexicon*, Cambridge, Mass.: MIT Press.

Pustejovsky, James, 1998, ‘The semantics of lexical underspecification’, Folia Linguistica 32, 323-347.

Reiter, R., 1980, 'Default logic', *Artificial Intelligence* 13, 81-132.

---

[10] Further details of this model are given in Hunter & Marten (1999).

Rips, Lance J., 1995, 'The current status of research on concept combination'. *Mind & Language* 10, 72-104.

Sag, Ivan A. & Anna Szabolcsi, eds., 1992, *Lexical Matters*, Stanford, Cal.: CSLI.

Sperber, Dan & Deirdre Wilson, 1986/1995, *Relevance: Communication and Cognition*, Oxford: Blackwell.

Sperber, Dan & Deirdre Wilson, 1997, 'The mapping between the mental and the public lexicon', *UCL Working Papers in Linguistics* 9, 107-125.

Wilson, Deirdre, 1999, 'Review Lecture of Lascarides, Copestake & Briscoe 1996 and Lascarides & Copestake 1998', UCL Pragmatics Reading Group.