Part I

Deliverable Identification Sheet

Project ref. no.	IST-2001-37836
Project acronym	DeepThought
Project full title	DeepThought — Hybrid Deep and Shallow Methods for Knowledge-Intensive Information Extraction

Security (distribution level)	Public
Contractual date of delivery	M06: April 2003
Actual date of delivery	April 11, 2003
Deliverable number	D1.1a
Deliverable name	Report on the design of RMRS (preliminary version)
Type	Report
Status & Version	Draft
Number of pages	20
WP contributing to the deliverable	WP1a
WP / Task responsible	U-CAM-CL
Other contributors	
Author(s)	Ann Copestake, U-CAM-CL
EC Project Officer	Yves Paternoster
Keywords	Multilingual Grammar Engineering
Abstract (for dissemination)	Deep and shallow processing techniques have complementary strengths and weaknesses. There are various ways in which deep and shallow techniques can be combined in practical systems to take advantage of the strengths of each approach. In this report, I argue for the advantages of using a semantic representation as the interface between various types of deep and shallow processing. I illustrate how with a flat semantics approach, shallow processing can be regarded as producing a semantics which is underspecified with respect to deeper processing, but fully compatible with it. I describe the notion of specificity which results and illustrate its connection with specificity in typed feature structures. I also show how semantic composition rules designed for deep formalisms can be adapted to shallow processing.

Contents

Part II

Executive summary

Deep and shallow processing techniques have complementary strengths and weaknesses. There are various ways in which deep and shallow techniques can be combined in practical systems to take advantage of the strengths of each approach. In this report, I argue for the advantages of using a semantic representation as the interface between various types of deep and shallow processing. I illustrate how with a flat semantics approach, shallow processing can be regarded as producing a semantics which is underspecified with respect to deeper processing, but fully compatible with it. I describe the notion of specificity which results and illustrate its connection with specificity in typed feature structures. I also show how semantic composition rules designed for deep formalisms can be adapted to shallow processing.

The introduction introduces the idea of combining deep and shallow processing and argues for the use of a semantic representation. The second section gives a detailed worked example as a way of introducing RMRS. The third section is an outline of the formalisation — this is currently incomplete and will be considerably expanded for the final version of this report. The fourth section introduces the use of RMRS with RASP and briefly describes the preliminary implementation of the software for creating RMRSs from RASP tree output. Section five goes into the underlying theory behind semantic composition for RMRS, describes why compatible composition for deep and shallow processing is important, and discusses how the existing semantic composition rules for the ERG/MRS are adapted to RASP/RMRS. The sixth section outlines some ways in which RMRS might be used, while the final section contains conclusions and an outline of further work.

Part III

Report on the design of RMRS (preliminary version)

Ann Copestake

aac@cl.cam.ac.uk

Abstract

Deep and shallow processing techniques have complementary strengths and weaknesses. There are various ways in which deep and shallow techniques can be combined in practical systems to take advantage of the strengths of each approach. In this report, I argue for the advantages of using a semantic representation as the interface between various types of deep and shallow processing. I illustrate how with a flat semantics approach, shallow processing can be regarded as producing a semantics which is underspecified with respect to deeper processing, but fully compatible with it. I describe the notion of specificity which results and illustrate its connection with specificity in typed feature structures. I also show how semantic composition rules designed for deep formalisms can be adapted to shallow processing. Some practical experiments are described using the tag sequence grammar from the RASP tools and the LinGO English Resource Grammar.[1]

1 Introduction

There are many highly-developed deep processing systems, often associated with linguistically-motivated formalisms including Head-driven phrase structure grammar (HPSG), Lexical-functional grammar (LFG), Tree-adjoining grammar (TAG, XTAG, LTAG), Generalized phrase structure grammar (GPSG) and various types of categorial grammar (CG). Some broad-coverage deep grammars have been developed, which cover a large proportion of linguistic phenomena. These grammars can be used to produce a detailed logical semantic representation. Formally, the grammars are expressed declaratively and in principle can be used for generation as well as analysis. Although formalisms such as HPSG have been considered to be practically intractable, well-engineered systems now exist which are fast enough for practical use in applications such as email response (see e.g., Oepen et al, 2003).

There are, however, at least three classes of problem which limit the practical utility of systems based on deep processing of text (speech raises other issues which will not be discussed here). The first is that the number of analyses returned can be far too high for any realistic prospect of filtering by a post-processor. In general, techniques such as packing merely delay the problem, since few applications can work off the packed representation. Integration of statistical techniques for parse selection with deep processing is promising, but is far from being a solved problem, especially with HPSG. The second problem is robustness: deep systems assume a notion of grammaticality so that input which does not meet the criteria embodied in the grammar fails to parse. For edited text, parse failures are generally caused by lack of coverage — often missing lexical information. Constructing detailed lexical information manually is time-consuming and error-prone, while automatic or semi-automatic techniques run into problems because of data sparsity. For unedited text, minor typographic mistakes can cause complete parse failure. The third issue is speed: deep processing techniques are not fast enough to process large volumes of text, as is required by information retrieval, information extraction and question answering, for instance.

There are a range of techniques for processing which provide some linguistic structure at a shallower level than the deep processors. These include part-of-speech tagging and NP chunking. Such techniques are much faster than deep processing and also more robust. Because they involve a less detailed representation, they tend to involve a lower degree of ambiguity, but more importantly, they are generally developed with a fully integrated

[1]This work has benefited greatly from discussions with Ted Briscoe, Dan Flickinger, Alex Lascarides, Stephan Oepen and Simone Teufel and with researchers at Saarbrücken and Edinburgh where versions of it were presented. As usual, however, all errors are the responsibility of the author.

statistical metric for selection. However, they do not construct sufficient information to support conventional semantic representation.

Parsers which produce relatively simple dependency structures without a detailed lexicon are intermediate in depth (e.g., the RASP system Briscoe and Carroll, 2002). They are faster and more robust than deep parsers, though not as efficient as the shallower techniques, but there is generally no treatment of long-distance dependencies, and the structures produced contain insufficient information to recover scope information. They allow ungrammatical strings and the lack of lexical subcategorization information still precludes conventional semantic representation.

Deep processing is inherently slower than shallow processing, because it involves a larger search space. It is inherently less robust, because it assumes some notion of grammaticality and requires more detailed lexical information. Conversely, although it is possible to enhance shallow processing to overcome some of the disadvantages mentioned above, this requires more search space and additional knowledge sources. Even if the additional knowledge can be automatically acquired, robustness will decrease, because of data sparsity. The deep/shallow dimension is somewhat associated with the distinction between manually- and automatically- constructed resources, but there are manually constructed shallow processing systems and some automatically learned intermediate processors. For the purposes of this report, the most important distinction between deep and shallow processing concerns the lexical requirements: in particular, deep processing requires detailed subcategorization information while shallow processing just needs some indication of possible parts-of-speech, plus information about irregular morphology.

Question answering is one application where there is a natural division of labour between deep and shallow processing. The requirement to process large volumes of answer text means that some shallow approach is required, at least for identification of likely answers. On the other hand, there are far fewer questions to process, the questions tend to be relatively short and precise analysis is important. So deep processing of questions is feasible and useful. In particular, questions can involve long-distance dependencies which will not be identified by shallow techniques.

Another application involves intelligent summarization. Key passages of text can be identified robustly using cue phrases. For instance, in scientific papers it is possible to identify sentences which describe the goal of the paper (Teufel, 1999). The vocabulary and syntactic structures used by authors to introduce the goal are relatively constrained and could potentially be deep parsed, although the actual text that describes the work performed is much less constrained and therefore difficult to deep parse with sufficient robustness.

Information extraction is also an application which should benefit from combining deep and shallow processing. An IE system needs to be able to process very large quantities of text, so it cannot be attempted by deep processing alone. On the other hand, most existing shallow processing systems rely on large amounts of hand-coding for each new domain, and are not very precise. The ideal situation would be to use shallow processing with limited domain-tuning to identify interesting sentences and then robustly invoke a deep processor with a domain-independent grammar to precisely extract the information which is of interest.

Deep processing may be sufficiently efficient for email answering. But lack of robustness causes problems, which could potentially be alleviated by using shallow processing as a backoff technique.

There are various strategies for combining deep and shallow processing (for specific approaches, see e.g., Crysmann et al, 2002; Daum et al, 2003). One involves augmenting a deep processor with shallow techniques. This may be appropriate in circumstances where deep processing is fast enough and the requirement is to increase robustness in the face of parse failure. For instance, a chart may be constructed via deep processing and, if full parsing fails, the substructures can be connected by more robust methods (e.g., Kasper et al, 1999). Another approach is to preprocess with a shallow parser and to use the information produced to guide the deeper parser. This is commonly done with POS tagging, which can be used as a preprocessor for deep parsing, but intermediate parsers could also be used in this way (if their output were compatible with the deeper processing). This cuts down ambiguity and thus increases speed, but only improves robustness if there is some strategy for backing-off to the shallow parse results if the deep parse fails. A less-explored strategy is to invoke shallow processing first and to deep process only selected parts of the input, as suggested above for IE.

All these strategies require integration of the deep and shallow processors, to differing degrees. Normally, integration at the syntactic level is attempted. For instance, a POS tagger's categories can be translated into the classes used by a deep parser in order to disambiguate. But even this case, which is the simplest we might consider, raises problems. For instance, a POS tagger may systematically tag *-ing* forms occurring in a prenominal positional as adjectives, while a deep parser might treat them as gerunds. Individual discrepancies like this do not pose a large problem, but each tagger/parser pair will require different adjustments and the cumulative effect is non-trivial. Much more serious issues arise with shallow parser that provide some bracketing of constituents, since the deep and shallow parsers may not agree on this, especially given that the underlying linguistic theory between different deep parsers gives rise to alternative bracketings. There seems little chance of agreement on a canonical syntactic structure. This is also illustrated by the difficulty of providing meaningful comparisons between parsers unless there is an agreed treebank. Carroll et al (1998) argue for comparison to be based on grammatical relations, but this also has its critics.

The alternative proposed here is to produce semantic representations from deep and shallow parsers, aiming for a maximal degree of semantic compatibility. The most important reason to take this approach is that work on providing a standardised semantic representation would be useful for many purposes: allowing different parsers (and generators) to be connected up to a variety of underlying systems. Standardising grammatical relations is not directly useful in this way. Furthermore, we can view semantic representations constructed from shallow parsers as underspecified forms of the output from a deep parser in a way that can be made formally precise. Finally, any constraints that arise from the domain are most naturally expressed in terms of semantics and it is potentially very useful to be able to apply domain constraints to both deep and shallow parsing.

The types of integration mentioned above are all possible under this strategy, on the assumption that semantic structures are produced compositionally. A deep parse failure involves disconnected semantic structures, which may be connected via shallow parsing, possibly with the help of domain constraints. Semantic structures constructed from a shallow parser can guide deep parsing, because only constituents with compatible semantics need be considered. Finally, in an application such as IE, the patterns necessary to identify regions of potential interest can be robustly and succinctly expressed in terms of semantics, and deep parsing can be used to further instantiate this information if shallow parsing is produces results which are too underspecified be be used by themselves.

To summarise, making shallow parsing return underspecified semantic representations has the following advantages:

- Integrated parsing: shallow parsed phrases could be incorporated into deep parsed structures.
- Deep parsing could be invoked incrementally in response to information needs.
- Reuse of knowledge sources is possible: information for further processing of shallow parsing might also be used on deep parser output. (e.g., domain knowledge, recognition of named entities, transfer rules in MT)
- The formal properties should be clearer, and thus the representations might be more generally usable.

In the remainder of this report, I will describe a proposed semantic representation that can function as the interface between deep and shallow processing. In section 2 I will give an initial simple example and illustrate how shallow parsers can be used to give compatible output to a deep parser. The specific semantic representation introduced here is a modification of Minimal Recursion Semantics (MRS: Copestake et al 1999), known as Robust MRS (RMRS). Section 3 gives more details of the formal properties of RMRS. Section 4 briefly describes RASP and introduces RASP-generated RMRSs. In section 5, I describe semantic composition with RMRS using rules associated with the RASP tag-sequence grammar, and show that it is compatible with the algebra used for MRS. Section 6 discusses some initial approaches for using RMRS, and finally section 7 contains conclusions and further work.

2 A simple example of RMRS

In this section, RMRS is introduced by example. I will consider the semantic representations for the sentence:

(1) every fat cat sat on some table

The scoped semantic representations for this sentence can be assumed to be the following:

$$\text{every}(x, \text{fat}(x) \wedge \text{cat1}(x), \text{some}(y, \text{table1}(y), \text{sit1}(e, x) \wedge \text{spast}(e) \wedge \text{on2}(e', e, y)))$$

$$\text{some}(y, \text{table1}(y), \text{every}(x, \text{fat}(x) \wedge \text{cat1}(x), \text{sit1}(e, x) \wedge \text{spast}(e) \wedge \text{on2}(e', e, y)))$$

In fact, this is a slightly simplified form of the scoped semantic representation output by the current version of the LinGO English Resource Grammar (ERG: Copestake and Flickinger, 2000). Here, I will assume that it is relatively uncontroversial. The representation uses generalized quantifiers — this is a fundamental assumption in what follows. It also assumes a neo-Davidsonian approach using event variables, and a simple account of tense as a predicate on events: these assumptions are less critical to the general RMRS idea. The assumption of the 'extra' event in prepositions is motivated by predicative use after a copula: this follows the current treatment in the ERG, but is not important for RMRS.

The central idea in this report is that this fairly conventional semantic representation can be directly related to very underspecified semantic information. To take the extreme case, I will argue that a POS tagger can be viewed as providing very underspecified semantic representation, and then show how this is formally related to the structures shown above.

2.1 Semantics from a POS tagger

A conventional POS-tagger output for the example sentence is as follows:

(2) every_AT1 fat_JJ cat_NN1 sat_VVD on_II some_DD table_NN1

The POS tagger can be viewed as providing some semantic information:

- The lemmatization identifies a class of relations. For instance, in the example above, *sat* is lemmatized to *sit*. On the usual assumption of a convention whereby relation names (i.e., predicates) are related to the spelling of the lemma, we can associate the tagged word with a name corresponding to the class of all relations that correspond to that spelling.
- The POS tag provides partial disambiguation of the class of relations corresponding to the lemma's spelling: e.g., we can distinguish the noun *table* from the verb.
- The POS tag for inflected forms can be used to identify the semantic contribution of the inflection.

However POS tagging provides no relational information and no information about scope. Intuitively, therefore, the semantic information content is very underspecified compared to the semantics from the deep processor.

The structure below is a proposal for the representation of the semantic information associated with the POS tagged sentence:

$$\text{every}(x0), \text{fat\_J}(x1), \text{cat\_N}(x2), \text{sit\_V}(e3), \text{past}(e3), \text{on}(e4), \text{some}(x5), \text{table\_N}(x6)$$

For now, this should be taken informally, as an indication of what the logical representation could contain. Details will be discussed later on. Note, however, that the representation identifies some entities which are involved, but essentially constructs one entity per word in the input: for instance, x0 and x1 are not related. The variables associated with the entities are sorted, so that 'x's are object-denoting and 'e's are event-denoting. Since there is no relational information, the predicates have to be unary. Notice also that the naming convention for the relations allows them to be generated automatically. Open class words have their broad part of speech indicated based on the tag.[2]

The crucial assumption is that this semantics can be generated from a 'lexicon' for the tags, rather than a genuine lexicon which is indexed by lexeme. Specifically, the tag lexicon assumed here is as follows:

AT1	lexrel(x)
JJ	lexrel_J(x)
NN1	lexrel_N(x)
VVD	lexrel_V(e), past(e)
II	lexrel(e)
DD	lexrel(x)

Here, 'lexrel' is a special symbol, which is taken to be replaced by the individual lemma. Producing the semantics from the tagger output and this lexicon is a trivial matter. All variables are taken to be different unless repeated in the same lexical entry. All tags correspond to one or more relation/argument combination (a relation plus its arguments will be referred to as an **elementary predication**). The elementary predications are simply concatenated to give the semantic representation for the sentence, hence no sort of tree structure is needed to guide composition.

At this point, there are two main questions:

[2] Our current assumption is that it is only necessary to do this for open class words, though obviously it could be done for closed class ones as well.

- Can the proposal that this is an underspecified semantic representation be made formally precise?
- How do we manipulate syntax of the semantic representation which is generated by deep processing in order to make the connection with the shallow semantics more obvious?

The answer to the first question follows established approaches to underspecification, in particular, Copestake et al (1999) and Copestake et al (2001). The shallow semantics is interpreted in terms of the full representations it can be expanded to. In what follows, I will go through the main steps required to answer the second question. The central idea is to split up the deep representation into minimal units which can be manipulated independently. The shallow semantics then contains only some of these units.

2.2 Flat semantics, underspecified scope (MRS)

The first step is to adopt a version of the deep semantics which is capable of underspecifying scope, and in which the elementary predications are flattened, so that the essential rule of composition is that elementary predications are appended. The approach I will assume is Minimal Recursion Semantics (MRS, Copestake et al, 1999), although Bos' (1995) semantics would be another candidate.

For convenience, the scoped structures for our sample sentence are repeated below:

$$\text{every}(x, \text{fat}(x) \wedge \text{cat1}(x), \text{some}(y, \text{table1}(y), \text{sit1}(e, x) \wedge \text{spast}(e) \wedge \text{on2}(e', e, y)))$$

$$\text{some}(y, \text{table1}(y), \text{every}(x, \text{fat}(x) \wedge \text{cat1}(x), \text{sit1}(e, x) \wedge \text{spast}(e) \wedge \text{on2}(e', e, y)))$$

The single MRS structure corresponding to these scopes is as follows:

$l0$:every$(x, h1, h2), l1$:fat$(x), l1$:cat1$(x), l4$:sit1$(e, x), l4$:spast$(e), l4$:on2$(e', e, y), l5$:some$(y, h6, h7), l6$:table1(y), qeq$(h1, l1)$, qeq$(h6, l6)$

I will not go through details of the MRS structure here, see Copestake et al (1999) for details. The salient points are:

- The MRS can be specialised to give exactly the two scoped structures.
- Each elementary predication is labelled $l0$, $l1$ etc.
- Scopal arguments in the scoped representation are indicated by holes: $h1$, $h2$ etc. Holes may be filled by labels according to the constraints on scope. In the case of the example above, the two possible scopes correspond to $h1 = l1, h2 = l5, h6 = l6, h7 = l4$ and $h1 = l4, h2 = l5, h6 = l6, h7 = l0$.
- Constraints on scope are given above by qeq conditions, but since the details of these are unimportant for the discussion in the current work, they will not be further explained here.

In standard MRS, implicit conjunction is indicated by equality between labels. For instance, l1 labels both $l1$:fat(x) and $l1$:cat1(x). However, this is not essential, and in what follows, I use a variant where labels on elementary predications are unique and conjunction arising from intersective modification is explicitly represented. This corresponds to:

$l0$:every$(x, h1, h2), l1$:fat$(x), l2$:cat1$(x), l3$:CONJ$, l4$:sit1$(e, x), l4$:spast$(e), l14$:on2(e', e, y),
$l9$:CONJ$, l5$:some$(y, h6, h7), l6$:table1(y),
qeq$(h1, l3)$, qeq$(h6, l6)$, in-g$(l3, l1)$, in-g$(l3, l2)$, in-g$(l9, l4)$, in-g$(l9, l14)$

Here CONJ stands for conjunction (uppercased because the relation is supplied by the construction and not lexically) and 'in-g' stands for 'in group': this representation allows for the fact that there may be an indefinitely large number of elements in a conjunction.

MRS takes us quite a long way towards the goal of splitting the semantics into individual units, but it still assumes that we have knowledge of the arity of the predicates, which does not apply to the shallow semantics situation. To allow for this, we need to move to a representation where arguments can be considered individually.

2.3 Parsons' representation

The following shows a more factorised representation where the base predicates are all maximally unary and the arguments are represented by separate (binary) relations:

$l0$:every(x), $l1$:fat(x), $l2$:cat1(x), $l3$:CONJ, $l4$:sit1(e), $l4$:spast(e), $l14$:on2(e'), $l9$:CONJ, $l5$:some(y), $l6$:table1(y),
qeq$(h1, l1)$, qeq$(h6, l6)$, in-g$(l3, l1)$, in-g$(l3, l2)$, in-g$(l9, l4)$, in-g$(l9, l14)$
RSTR$(l0, h1)$, BODY$(l0, h2)$, RSTR$(l5, h6)$, BODY$(l5, h7)$,
ARG1$(l4, x)$, ARG1$(l14, e)$, ARG2$(l14, y)$

Conversion from MRS to this representation merely involves some naming convention for the arguments of each type of predicate. The default naming convention is ARG1, ARG2 etc, with some relations corresponding to closed class words, such as determiners, having more characteristic argument names (e.g. RSTR and BODY for quantifiers). In this respect, the notation is close to the typed feature structures used internally to represent semantics in HPSGs such as the ERG.

I refer to this as 'Parsons style' notation, since it relates to Parsons' (1990) proposal for making arguments first-class predications. For instance, Parsons argued that kick(e) $\wedge$ agent(e, x) $\wedge$ patient(e, y) is a preferable notation to kick(e, x, y). Parsons reasons for advocating this approach are somewhat irrelevant to our concerns in this report, but the reason for using this style of notation here is directly related to Dowty's (1989) discussion of Parsons' approach. Dowty observed that this style of representation is better suited to nominalizations than the conventional logical form because nominals' arguments are generally optional, and this notation makes it easier to omit arguments. The lack of lexicon in shallow semantics means we can't know the arity of open-class predicate, but Parsons' style of representation allows incremental addition of arguments. This amounts to treating them as optional in terms of the syntax of the semantics.

There are two significant differences from Parsons. The first is the use of ARG1 etc rather than agent/patient. The primary reason for this is to allow a simple scheme which does not require a lexicon: it is possible to map it into a more detailed scheme (this is discussed in Copestake and Flickinger, 2003). As I will discuss in §3.1, underspecification of ARG-relations allows for dependency structures where we know there is some relation but can't be sure which one it is.

The second distinction is the use of labels instead of events as the anchor for the arguments. This allows the same approach to be used for cases where events are not motivated such as relational nouns (e.g., *group*) and scopal adverbs and connectives (e.g., *not*, *probably*, *but*).

This notation requires that individual elementary predications have distinct labels. This is why the implicit representation of conjunction in 'standard' MRS is not suitable and the alternative representation was adopted above.

2.4 Distinct variable names, plus equalities

The separation of arguments is the major difference between the proposed representation and MRS, but there are two further steps. The first is to have distinct names for variables and to represent identity between them with equalities. The reason for this is to make it clear what extra information is involved: the information that variables are equated can be added incrementally.

$l0$:every$(x0)$, $l1$:fat$(x1)$, $l2$:cat1$(x2)$, $l3$:CONJ, $l4$:sit1$(e3)$, $l4$:spast$(e3)$, $l14$:on2$(e4)$, $l9$:CONJ,
$l5$:some$(x5)$, $l6$:table1$(x6)$,
qeq$(h1, l1)$, qeq$(h6, l6)$, in-g$(l3, l1)$, in-g$(l3, l2)$, in-g$(l9, l4)$, in-g$(l9, l14)$
RSTR$(l0, h1)$, BODY$(l0, h2)$, RSTR$(l5, h6)$, BODY$(l5, h7)$ARG1$(l4, x2)$, ARG1$(l14, e3)$, ARG2$(l14, x5)$,
$x0 = x1, x1 = x2, x5 = x6$

2.5 A predicate naming convention

Finally, a predicate naming convention is required so that the automatically constructed predicates in the tagger output can be straightforwardly related to the deep representation. This is indicated below:[3]

$l0$:every$(x0), l1$:fat_J$(x1), l10$:cat_N$(x2), l3$:CONJ$, l4$:sit_V_1$(e3), l4$:spast$(e3), l14$:on$(e4), l9$:CONJ,
$l5$:some$(x5), l6$:table_N_1$(x6)$,
qeq$(h1, l1)$, qeq$(h6, l6)$, in-g$(l3, l1)$, in-g$(l3, l2)$, in-g$(l9, l4)$, in-g$(l9, l14)$
RSTR$(l0, h1)$, BODY$(l0, h2)$, RSTR$(l5, h6)$, BODY$(l5, h7)$, ARG1$(l4, x2)$, ARG1$(l14, e3)$, ARG2$(l14, x5)$,
$x0 = x1, x1 = x2, x5 = x6$

Here sit_V_1 is more specific than sit_V, spast more specific than past, and table_N_1 more specific than table_N. This is discussed further in §3.1.

2.6 Robust underspecification

We have now arrived at a point where it can easily be seen that the tagger output is an underspecified form of the deep representation. The tagger output was:

every$(x0)$, fat_J$(x1)$, cat_N$(x2)$, sit_V$(e3)$, past$(e3)$, on$(e4)$, some$(x5)$, table_N$(x6)$

The version with labels is:

$l0$:every$(x0), l1$:fat_J$(x1), l10$:cat_N$(x2), l4$:sit_V$(e3), l4$:past$(e3), l14$:on$(e4), l5$:some$(x5), l6$:table_N$(x6)$

This can be converted into the deep representation shown above by adding 'CONJ's, 'in-g's, 'qeq's, ARGn relations and equalities and by specialising predicates. To make the relationship clearer, I have used the same variable names in both cases, but in general, alphabetic variance has to be taken into account.

Notice that the output from an NP chunker is intermediate in specificity, for instance:

$l0$:every$(x0), l1$:fat_J$(x1), l10$:cat_N$(x2), l4$:sit_V$(e3), l4$:past$(e3), l14$:on$(e4), l5$:some$(x5), l6$:table_N$(x6)$,
$l3$:CONJ, in-g$(l3, l1)$, in-g$(l3, l2)$, $x0 = x1, x1 = x2$

The extraction of semantic representations which are of intermediate specificity is discussed in some detail in §4 with reference to the RASP system.

2.7 Lexical requirements for RMRS

In general terms, the lexical requirements for construction of RMRSs are as follows:

1. A tag lexicon or similar class-based lexicon is needed for shallow processing.
2. Consistent naming conventions are required for relations corresponding to open-class words, depending only on the orthography of the lemma and on the POS tag.
3. The representations produced by deeper analysis must specify their relationship to shallower predicates if they make finer-grained distinctions:
   e.g., sit_V_1 must be stated to be more specific than sit_V
4. There is a consistent convention for ARG1 etc which depends on syntax.
5. A lexicon is required for some closed class words for shallow/intermediate processing (e.g., we don't want any semantics for some closed class words such as *do*). On the assumption that the semantics from the deep analsys is taken as normative, this forms part of the meta-level semantic interface (SEM-I) which is discussed in Copestake and Flickinger (2003).
6. No lexicon is required for open-class words.

[3]The examples in this report should be taken as indicative rather than normative: the detailed predicate naming convention remains to be decided.

2.8 Indexing

The problem of comparing two RMRS structures is potentially exponential. In the particular case where the structures have been built from the same string, indexing can be used to allow for efficient comparison.[4] We propose to index by characters, with the origin of the elementary predications being shown by a character range. For instance, consider the example input sentence:

(3) Every cat sat on some table.

The characters can be indexed as follows:

e	v	e	r	y		f	a	t		c	a	t		s	a	t
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
	o	n		s	o	m	e		t	a	b	l	e	.		
17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		

Indexing the elementary predications leads to the following tagger output, where the indices are shown before the elementary predications:

$\langle 0,4\rangle$:$l0$:every$(x0)$, $\langle 6,8\rangle$:$l1$:fat_J$(x1)$, $\langle 10,12\rangle$:$l10$:cat_N$(x2)$, $\langle 14,16\rangle$:$l4$:sit_V$(e3)$, $\langle 14,16\rangle$:$l4$:past$(e3)$,
$\langle 18,19\rangle$:$l14$:on$(e4)$, $\langle 21,24\rangle$:$l5$:some$(x5)$, $\langle 26,30\rangle$:$l6$:table_N$(x6)$

The reason for using character positions rather than word positions to index is that we cannot assume that two systems will tokenize the input in exactly the same way. The character position indices must be assigned on the basis of the raw input, before any sort of preprocessing has been carried out.

A complication arises when elementary predications are introduced by a construction. In this case, they are indexed by the initial and final character string of the elements of the construction (see CONJ in the example below).

$\langle 0,4\rangle$:$l0$:every$(x0)$, $\langle 6,8\rangle$:$l1$:fat_J$(x1)$, $\langle 10,12\rangle$:$l10$:cat_N$(x2)$, $\langle 14,16\rangle$:$l4$:sit_V$(e3)$, $\langle 14,16\rangle$:$l4$:past$(e3)$,
$\langle 18,19\rangle$:$l14$:on$(e4)$, $\langle 21,24\rangle$:$l5$:some$(x5)$, $\langle 26,30\rangle$:$l6$:table_N$(x6)$,
$\langle 6,12\rangle$:$l3$:CONJ, in-g$(l3,l1)$, in-g$(l3,l2)$, $x0=x1, x1=x2$

Since qeqs, equalities etc relate variables which are introduced by elementary predications, there should be no need to index these. The variable equivalences across two RMRSs are established by the equivalences between elementary predications.

3 An outline of the approach to specificity

Specificity is defined in terms of the syntax of the semantic representation, not the denotation (cf MRS algebra semantics, Copestake et al, 2001, where the models are fully specified LFs).

Informal definition: a semantic representation A is an underspecified form of B if A can be converted to B by:

1. adding argument-relations (e.g., ARG1)
2. adding equalities between variables
3. specialising predicates
4. adding non-lexical elementary predications (e.g., CONJ)

The formal properties depend on treating representations as made up of minimal units of semantic information. Specificity is then treated in terms of a partial order on structures in terms of information containment (cf (typed) feature structures). This will be described in more detail in a later version of this document.

[4] Believed to be linear time, but this remains to be confirmed.

3.1 Underspecification of relations

The partial order on individual elementary predications is determined by a partial order on predicates. For instance: sit_V1 $\sqsubset$ sit_V, spast $\sqsubset$ past

There is also a partial order on the argument relations, so that ARG1 $\sqsubset$ ARGN etc. This allows for the situation where an analyser can determine that some argument position is filled by a variable, but the specific position cannot be determined. A refinement of this is to use a hierarchy of relations:
ARG2 $\sqsubset$ ARG2-4, ARG3 $\sqsubset$ ARG2-4, ARG4 $\sqsubset$ ARG2-4, ARG2-4 $\sqsubset$ ARGN etc.

For further details of the relationship to thematic roles, see Copestake and Flickinger (2003).

4 Intermediate level processing

The interconversion of MRSs produced by the ERG and RMRSs is largely straightforward, assuming that the argument relation naming convention is explicitly known. In this section, I consider constructing RMRSs from the RASP system (Robust accurate domain-independent statistical parsing: Briscoe and Carroll, 2002). RASP consists of a tokenizer, POS tagger, lemmatizer, tag-sequence grammar, statistical disambiguator. The central part of RASP from our current perspective is the grammar, which is a detailed grammar for (written) English, written in a constraint-based formalism. RASP is robust because the lexicon only contains POS tags. For instance:

```
WORD VVD : V0[FIN +, AUX -, VFORM PAST, CONJ -].
```

The rules mention SUBCAT, but this is not lexically instantiated:

```
PSRULE S/np_vp : V2[WH -, INV -] --> N2 H1[FIN +].
PSRULE V/np : V1 --> H0[VSUBCAT NP] N2.
```

The system is in fact used for experiments on acquisition of subcategorization information. It can be regarded as being of intermediate depth: it is quite robust, but uses a weak notion of grammaticality and has minimal lexical requirements. It supports a variety of output formats, including grammatical relations and parse trees. The current work extends this by allowing RMRS output.

4.1 Robust semantics from RASP

The current system for constructing RMRSs from RASP works off the tree output, keyed by the grammar rules. For instance:

```
(|T/txt-sc1/----|
 (|S/np_vp| (|NP/det_n1| |Every:1_AT1| (|N1/n| |cat:2_NN1|))
  (|V1/v| |bark+ed:3_VVD|)))
```

gives rise to the following RMRS:
$h1$:_every_rel$(x2)$, $h3$:_cat_N_rel$(x2)$, $h5$:_bark_V_rel$(e6)$, ARG1$(h5, x2)$
This is actually output in XML format:[5]

```
<rmrs>
<ep><realpred>_every_rel</realpred><label>h1</label><arg>x2</arg></ep>
<ep><realpred>_cat_N_rel</realpred><label>h3</label><arg>x2</arg></ep>
<ep><realpred>_bark_V_rel</realpred><label>h5</label><arg>e6</arg></ep>
<rarg><rargname>ARG1</rargname><label>h5</label><arg>x2</arg></rarg>
</rmrs>
```

In fact, a more detailed structure could be produced including qeq relations for scope, but I will ignore this complication for now.

The tag 'lexicon' that is used in the composition is the following:

[5] The current DTD is given in the appendix but the details of the XML format will be finalised in a later version of this document.

```
<le>
<tag>AT1</tag>
<semstruct>
<hook><index>X</index><label>H</label></hook>
<ep><pred></pred><label>H</label><arg>X</arg></ep>
</semstruct>
</le>

<le>
<tag>NN1</tag>
<semstruct>
<hook><index>X</index><label>H</label></hook>
<ep><pred></pred><label>H</label><arg>X</arg></ep>
</semstruct>
</le>


<le>
<tag>VVD</tag>
<semstruct>
<hook><index>E</index><label>H</label></hook>
<ep><pred></pred><label>H</label><arg>E</arg></ep>
</semstruct>
</le>
```

The lexical entries identify a ‘hook’ (see next section) and one or more elementary predications.

The grammar rules which are keyed off the tree are:

```
<rule>
<name>T/txt-sc1</name>
<dtrs>S OPT OPT OPT OPT</dtrs>
<head>S</head>
</rule>

<rule>
<name>S/np_vp</name>
<dtrs>NP VP</dtrs>
<head>VP</head>
<semstruct>
<hook><index>E</index><label>H</label></hook>
<rarg><rargname>ARG1</rargname><label>H</label><arg>X</arg></rarg>
</semstruct>
<eq>X NP.INDEX</eq>
<eq>H VP.LABEL</eq>
</rule>

<rule>
<name>NP/det_n1</name>
<dtrs>DET N</dtrs>
<head>N</head>
<eq>N.INDEX DET.INDEX</eq>
</rule>

<rule>
<name>N1/n</name>
```

```
<dtrs>N</dtrs>
<head>N</head>
</rule>

<rule>
<name>V1/v</name>
<dtrs>V</dtrs>
<head>V</head>
</rule>
```

Grammar rules identify the semantic head daughter (which is not necessarily identical to the syntactic head). They may also specify argument relations and elementary predications. This approach will be described in more detail in a later document. The formalism for RMRS composition in general is discussed in more detail in the next section.

5 An outline of RMRS composition

Semantic composition needs to be formalised for both deep and shallow systems. Since we are interested in investigating tight integration of deep and shallow processing, where we may need to combine phrases produced by a deep processor with those from a shallower system, the issue of composition is very important. In particular, compatible composition rules ensure that this sort of integration is feasible.

In this section, I briefly introduce the MRS composition principles described in Copestake et al (2001) and go on to show their relevance to composition with a POS tagger (trivially) and RASP.

5.1 A semantic algebra

Copestake et al (2001) define a semantic algebra which is intended to capture 'best practice' in semantic composition in an HPSG using MRS. Although semantic composition is implemented in terms of operations on typed feature structures, it is desirable to restrict the way that semantic structures are manipulated, both as a way of capturing linguistic generalizations about composition and also as an aid to designing efficient algorithms. The algebra is illustrated below with a simple example (ignoring scope). Semantic structures consist of four basic parts:

hook Hooks are used to give arguments for relations. They can be thought of as pointers into the relations which are set up lexically: they are the only way of accessing parts of the semantics of a sign. Hooks are functionally similar to lambda variables, although formally quite different. In a full system, hooks consist of three parts: an index, a label (the local top in MRS terms, see Copestake et al, 1999), and an external argument.

slots Slots are structures corresponding to syntactic unsaturation which also specify how the semantics is combined. A syntax 'slot' in one sign will be instantiated by being equated with a hook in another sign.

rels The bag of elementary predications. This is always accumulated in composition, though some rules may add relations themselves.

equalities The results of equating slots and hooks.

For scoped structures, qeqs are also required, but these are not important here.

Composition always involves the following principles:

1. the rels of the phrase are equal to the append of the the rels of the daughters plus any rels contributed by the construction
2. the equalities of the phrase are equal to the append of the the equalities of the daughters plus any equalities contributed by the construction
3. a (syntactically specified) slot in one structure (the daughter which corresponds to the **semantic head**) is filled by the hook of the other structure (by adding equalities)
4. the hook of the phrase is the semantic head's hook

5. the slots of the phrase are the unfilled slots of the daughters

The access to the semantics of a sign during composition is only via the hook and the slots.

The following should serve to illustrate the algebra (ignoring scope, qeqs etc):

	$hook$	$slots$	$rels$	$equalities$
cat :	$[x_2]$	$\{\}$	$[\text{cat}(x_2)]$	$\{\}$
every :	$[x_3]$	$\{[x_3]_{spec}\}$	$[\text{every}(x_3)]$	$\{\}$
every cat :	$[x_3]$	$\{\}$	$[\text{every}(x_3), \text{cat}(x_2)]$	$\{x_3 = x_2\}$

Here, the hook of *cat* has instantiated the spec slot of *every*, which is the semantic head (though not the syntactic head in the ERG). This leads to the equality between the variables in the result. The relations are accumulated. Since the slot is saturated, it is not carried on to the phrase.

The hook of *every cat* can then instantiate the subject slot of *sleep*, as below:

every cat :	$[x_3]$	$\{\}$	$[\text{every}(x_3), \text{cat}(x_2)]$	$\{x_3 = x_2\}$
sleeps :	$[e_1]$	$\{[x_1]_{subj}\}$	$[\text{sleep}(e_1, x_1)]$	$\{\}$
every cat sleeps :	$[e_1]$	$\{\}$	$[\text{sleep}(e_1, x_1), \text{every}(x_3), \text{cat}(x_2)]$	$\{x_1 = x_2, x_3 = x_2\}$

5.2 Semantic composition with a POS tagger

The semantic composition rules apply trivially in the case of POS-tagged text. In this situation, we have no knowledge of any syntax above the POS tag level, so there are no slots and the hooks are irrelevant. The composition principle of accumulation of elementary predications holds, however, so the semantics of the result involves an accumulation of the rels as specified in principle 1 above.

5.3 Simple semantic composition with RASP

The semantic composition mechanism in RASP obeys the principles of composition given above. It differs from the mechanism for composition in MRS in the ERG in some ways:

- Since there is no lexical subcategorization information in RASP, slots are always specified by rules rather than lexically.
- Rules have their own semantic structures.
- The algebra formalisation requires binary combination: in order to model this, we can assume that rules are combined with daughters one-by-one. However, in practice, the combination of a single rule with all the daughters can be carried out in a single step. Generally one of the daughters is the semantic head, but sometimes the rule itself is, in effect, the semantic head.
- If no semantic rule is specified corresponding to a rule used in a tree, the rels are appended. If there is more than one daughter, the hook of the result contains new variables, since any of the daughters could correspond to the semantic head. Semantic composition is thus robust to omissions in the semantic component of the grammar: if no semantics is associated with any rules in the tree, the result is equivalent to a POS-tagged RMRS.

Despite these differences, the semantic structures that are output from phrases in RASP are, in principle, compatible with those that would be output by the ERG.

The structures below should illustrate this. The slots are indexed by the daughter names in the rule. The hooks and slots are shown with labels included, since this is needed for the ARG1 relation:

	$hook$	$slots$	$rels$	$equalities$
every cat :	$[l_8, x_3]$	$\{\}$	$[\text{every}(x_3), \text{cat}(x_2)]$	$\{x_3 = x_2\}$
sleeps :	$[l_9, e_1]$	$\{\}$	$[\text{sleep}(e_1)]$	$\{\}$
S/np_vp :	$[l_3, e_4]$	$\{[l_9, x_8]_{NP}, [l_3, e_4]_{VP}\}$	$[\text{ARG1}(l_3, x_2)]\{\}$	
	$[l_3, e_4]$	$\{\}$	$[\text{sleep}(e_1), \text{every}(x_3), \text{cat}(x_2), \text{ARG1}(l_3, x_2)]$	$\{x_3 = x_2\}$

6 Using RMRS

This section contains a range of tentative proposals for methods of integrating deep and shallow processing via RMRS.

6.1 Question answering

A quite straightforward use of RMRS which has been tried in a preliminary fashion is as part of a question-answering system.[6] The answer texts were shallow parsed with the RASP system, partial RMRSs were extracted, and compared with RMRSs extracted from the questions via the LinGO ERG. This involves a robust, weighted comparison between the RMRSs.

6.2 Incremental invocation of deep parsing on identified chunks of text

In some cases, we can identify 'interesting' subparts of text by means of specific cue phrase. This was mentioned in the introduction in the context of Teufel's work, but the example used here concerns IE. For instance, suppose as part of an IE system, we are looking for comparative judgements about cell phones. The following sentence can be identified as potentially relevant, because it contains the cue *I prefer X to Y*:

(4) I really prefer the Nokia to the Motorola because it's about 20 grams lighter.

While it is implausible that we can fully deep parse all the sentences in an IE application, because of parsing speed and the work needed to tune to each new domain, it is reasonable that we should be able to deep parse the cues, since these will be of restricted length and less domain dependent. For instance, *prefer* is likely to be used in a similar way whether talking about cell phones, motorcycles or wine. It is thus plausible that we can manually create a suitable lexicon for the cue texts. The reason for deep parsing is that we want to be able to identify arguments more precisely than can be done with shallow parsing alone (for instance, in the example above, we want to be able to utilise information about the subcategorization of *prefer*) and because we want to avoid false positives (for instance, where an adverb conveying negation is used).

The proposed methodology is to construct a library of cues expressed using RMRS, and to use this library to select regions of shallow-parsed text for further deep processing. For instance, in the example above, the cue phrase could be expressed as follows:

$$l0\text{:prefer\_V}(e), \text{ARG1}(l0, x1), \text{ARG2}(l0, x2), \text{ARG3}(l0, x3), l1\text{:pron}_{1st}(x1) \rightarrow \text{BETTER}(x2), \text{WORSE}(x3)$$

The left hand side of the rule will match some piece of text, while the right hand side is intended to be indicative of some action that might result from recognition of the cue, such as populating some database. These sort of rules can be regarded either as a form of pattern-matching on semantic representations, or as a variety of shallow inference: the 'munging' rule machinery in the LKB currently supports something similar on MRSs.

A shallow parser that output an RMRS structure would serve to identify a possible cue. For instance, assume that a POS tagger plus NP chunker was being used on the example above. The left hand side of the cue would unify with the RMRS extracted by the shallow parser, so matching it. The deep parser can then be invoked on the textual region identified by the cue. The interface between the shallow and deep parser is the original text aligned with the indexed RMRS, so the cue can be associated with a specific region of text (possibly non-contiguous, if parenthetical information is excluded). The deep parser would confirm that the cue matched and would identify the two relevant arguments to *prefer*, enabling the database to be populated.

The most plausible way of building the RMRS cue library is to hand-annotate some sentences to identify cue phrases, deep parse the cues to produce RMRSs (with manual checking/parse selection) and then to have a semi-automatic method of generalising over sets of cues.

6.3 Patching up failed deep parses

The compatibility of composition between deep and shallow processing also makes RMRS a promising approach for connecting fragmentary deep parses. The idea is to take the phrases that can be deep analysed, and connect

[6]These experiments have been done as part of the language practicals in the Masters course in Computer Speech Text and Internet Technology, run by Ted Briscoe, Simone Teufel and the author.

them via the semantics given by the shallower analyser, using structures with compatible bracketing. For instance, suppose the deep analyser has produced an analysis of the following sentence with the bracketed components analysed, but without a complete parse being found, perhaps because the grammar does not allow adverbs before subcategorized arguments:

(5) (The American consultant)$_{np}$ (expected)$_{v}$ (us)$_{np}$ (unknowingly and inadvertently)$_{adv}$ (to compromise the integrity of our software)$_{vp}$.

A shallow parser that had fewer constraints might analyse this sentence, attaching the *to* phrase to the verb with an ARGN relationship. This analysis would be compatible with the information from the structures produced by the deep parser, allowing those to be connected. In fact, given reasonable semantic constraints on the relation for *expect*, the ARGN could be specialised to ARG3.

It is currently unclear whether the semantics output by, say, the RASP parser could be made sufficiently tightly compatible with the ERG to allow this approach to work. This approach would probably require the semantics for closed-class words to be largely consistent between RASP and the ERG, which requires that part of the semantic interface to the ERG (the SEM-I: Copestake and Flickinger, 2003) be in place.

6.4 Unknown words/phrases

RMRS might also be useful as part of an alternative to an unknown word/phrase mechanism for deep parsing. If a shallow parser can identify a phrase as a particular sort of named entity, for example, the deep parser can incorporate a structure corresponding to that phrase in the parse. For instance, consider the following example:

(6) Through howling wind and freezing rain the tourists keep coming to the place where Johnny 'Mad Dog' Adair used to live.

Assume that the deep parser fails on *Johnny 'Mad Dog' Adair* but that a named-entity recogniser can analyse it. A basic template for an NP allows a phrasal sign to be constructed which can be syntactically incorporated into the parse: if the named entity recogniser can produce RMRS structures, then the semantics will also be coherent. The main question for the approach is whether the RMRS that is produced by the named entity recogniser should be equivalent to something that would be constructed by the deep parser if the phrase could be analysed (tight integration) or simply a structure that gives rise to a well-formed semantic result, potentially with different internal information than would arise from the deep parsing process, if that were to succeed. For example, as part of an unknown word mechanism, a named entity system might produce an RMRS such as:

$$[h0, x][h0\text{:def}(x, h1, h2), h4\text{:person-name}(x), \text{NAME}(h4, `JohnnyAdair'), \text{NICKNAME}(h4, `MadDog')]$$

The use of the quantifier means that this is structurally compatible with the MRS that would be produced by the ERG, but the ERG could not easily be modified to construct a comparable internal structure. It remains to be seen whether tight integration is worth pursuing in this case.

Named entity recognition would also be a useful refinement to the idea of cue phrase matching, described above. Some regions of the text passed to the deep-parser could be marked-up by a named-entity recogniser associated with shallow parsing. The deep parser then does not need to access the internal structure of these phrases, but can instead simply incorporate them into the result.

6.5 Parser evaluation

Briscoe et al (2002) suggest that RMRS might be useful as an output format for parser evaluation. Since RMRSs are maximally factorised, calculating the degree of overlap between a parser output and some gold standard is possible. For this to be viable, there needs to be agreement about the gold standard account of a wide range of phenomena, so eventually this would require something along the lines of the SEM-I described in Copestake and Flickinger (2003).

7 Conclusions and further work

The main idea proposed here is that we can manipulate the syntax of the semantic representation so it can be split into minimal composition units suitable for a variety of processors. By factoring the semantics into minimal units, we have a basis for comparison between systems that can produce differing amounts of semantic information. The formalism allows for radical underspecification of semantic content. So far this looks promising, since many ideas seem to flow from this, but considerably more work needed to see how plausible it really is in practical systems.

The current status of RMRS software support is as follows:

- A preliminary XML syntax has been defined (see appendix) as has a more readable syntax.
- An initial version of an MRS-to-RMRS converter has been constructed for the LKB which works on the ERG and should in principle be adequate for other similar grammars.
- RASP-to-RMRS software has been built with a preliminary set of rules for the current tag sequence grammar (TSG 14).

The next steps are:

- Develop the subsumption test code.
- Improve the TSG14 semantic rules and test RASP/ERG compatibility of the semantic test suite.
- Develop the first version of the SEM-I to allow tighter ERG/RASP integration.
- Improve the LKB munging rule mechanism so it can be used with RMRS as discussed above.

7.1 Further work

One extension to the RMRS formalism which is likely to be necessary is the addition of inequalities. For instance, consider the following example:

(7) the donkey kicked the cat

This might receive the following semantics:

(8) $l0$:the$(x1), l1$:donkey$(x1), l2$:kick$(e2), l3$:the$(x3), l4$:cat$(x3)$, ARG1$(l2, x1)$, ARG2$(l2, x3)$

We don't want to be able to further specialise to:

(9) $l0$:the$(x1), l1$:donkey$(x1), l2$:kick$(e2), l3$:the$(x3), l4$:cat$(x3)$, ARG1$(l2, x1)$, ARG2$(l2, x3), x1 = x3$

Furthermore, 8 shouldn't be compatible with:

(10) $l0$:the$(x1), l1$:donkey$(x1), l2$:kick$(e2), l3$:the$(x3), l4$:cat$(x3)$, ARG1$(l2, x3)$

The most likely solution to this problem is to introduce inequalities:

(11) the donkey kicked the cat

(12) $l0$:the$(x1), l1$:donkey$(x1), l2$:kick$(e2), l3$:the$(x3), l4$:cat$(x3)$, ARG1$(l2, x1)$, ARG2$(l2, x3), x1 \neq x3$

It is important to note that we are not talking about the real world identity of entities, but rather about the models to which a given RMRS can be specialised. It seems reasonable to assume inequalities between all nominal and verbal 'base' variables. One possibility is to assume inequalities by default and allow equalities to override them. I should note in this context that the formalisation of RMRS in terms of partial order of information content should be essentially compatible with the version of defaults described by Lascarides and Copestake (1999), although this suggestion remains to be investigated in detail.

Appendix

```
<!ELEMENT rmrs (ep|rarg|ing|hcons)*>
<!ELEMENT ep (realpred, label, arg)>
<!ELEMENT rarg (rargname, label, arg)>
<!ELEMENT realpred (#PCDATA)>
<!ELEMENT rargname (#PCDATA)>
<!ELEMENT arg (#PCDATA)>
<!ELEMENT label (#PCDATA)>
<!ELEMENT ing (ing-a, ing-b)>
<!ELEMENT ing-a (#PCDATA)>
<!ELEMENT ing-b (#PCDATA)>
<!ELEMENT hcons (hreln, hi, lo)>
<!ELEMENT hreln (#PCDATA)>
<!ELEMENT hi (#PCDATA)>
<!ELEMENT lo (#PCDATA)>
```

An example:

```
<?xml version='1.0'?> <!DOCTYPE rmrs SYSTEM
"/homes/aac10/lingo/newlkb/src/rmrs/rmrs.dtd" >
<rmrs>
<ep><realpred>_bark_V_rel</realpred><label>h9</label><arg>e2</arg></ep>
<ep><realpred>_dog_N_rel</realpred><label>h8</label><arg>x5</arg></ep>
<ep><realpred>_def_rel</realpred><label>h4</label><arg>x5</arg></ep>
<ep><realpred>prpstn_rel</realpred><label>h1</label><arg>h3</arg></ep>
<rarg><rargname>ARG1</rargname><label>h9</label><arg>x5</arg></rarg>
<rarg><rargname>RSTR</rargname><label>h4</label><arg>h6</arg></rarg>
<rarg><rargname>BODY</rargname><label>h4</label><arg>h7</arg></rarg>
<hcons><hreln>QEQ</hreln><hi>h3</hi><lo>h9</lo></hcons>
<hcons><hreln>QEQ</hreln><hi>h6</hi><lo>h8</lo></hcons>
</rmrs>
```

Part IV

Bibliography and references

Bos, Johan (1995) 'Predicate Logic Unplugged', *Proceedings of the Tenth Amsterdam Colloquium,* Amsterdam, Holland.

Briscoe, Ted, John Carroll, Jonathan Graham, Ann Copestake (2002) 'Relational Evaluation Schemes', *Proceedings of the Beyond PARSEVAL Workshop at LREC 2002,* Las Palmas, Gran Canaria, pp. 4–8.

Briscoe, Ted and John Carroll (2002) 'Robust accurate statistical annotation of general text', *Proceedings of the 3rd International Conference on Language Resources and Evaluation (LREC-2002),* Las Palmas, Gran Canaria.

Carroll, John, Ted Briscoe, and Antonio Sanfilippo (1998) 'Parser evaluation: a survey and a new proposal', *Proceedings of the 1st International Conference on Language Resources and Evaluation (LREC-1998),* Granada, Spain, pp. 447–454.

Copestake, Ann, Dan Flickinger, Ivan Sag and Carl Pollard (1999) 'Minimal Recursion Semantics: An introduction', ms., CSLI.

Copestake, Ann and Dan Flickinger (2000) 'An open-source grammar development environment and broad-coverage English grammar using HPSG', *Proceedings of the Second conference on Language Resources and Evaluation (LREC-2000),* Athens, Greece, pp. 591–600.

Copestake, Ann, Alex Lascarides and Dan Flickinger (2001) 'An Algebra for Semantic Construction in Constraint-based Grammars', *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics (ACL 2001),* Toulouse, France.

Copestake, Ann and Dan Flickinger (2003) 'The semantic interface to the ERG and other LinGO grammars', Deep Thought website.

Crysmann, Berthold, Anette Frank, Bernd Kiefer, Stefan Mueller, Guenter Neumann, Jakub Piskorski, Ulrich Schaefer, Melanie Siegel, Hans Uszkoreit, Feiyu Xu, Markus Becker and Hans-Ulrich Krieger (2002) 'An integrated architecture for shallow and deep processing', *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL 2002),* University of Pennsylvania, Philadelphia, USA, pp. 441–449.

Daum, Michael, Kilian Foth and Wolfgang Menzel (2003) 'Constraint Based Integration of Deep and Shallow Processing Techniques', *Proceedings of the 10th Meeting of the European Chapter of the Association for Computational Linguistics (EACL-2003),* Budapest, Hungary.

Dowty, David (1989) 'On the semantic content of the notion 'thematic role'' in Gennaro Chierchia, Barbara Partee and Ray Turner (ed.), *Property Theory, Type Theory and Natural Language Semantics,* Reidel, Dordrecht, The Netherlands, pp. 69–129.

Kasper, W., Kiefer, B., Krieger, H., Rupp, C., Worm, K. (1999) 'Charting the depths of robust speech parsing', *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (ACL '99),* College Park, Maryland, USA.

Lascarides, Alex and Ann Copestake (1999) 'Default representation in constraint-based frameworks', *Computational Linguistics,* **25.1**, 55–106.

Stephan Oepen, Daniel Flickinger, Hans Uszkoreit and J-I. Tsujii (2003) *Collaborative Language Engineering: A Case Study in Efficient Grammar-based Processing,* CSLI Publications.

Parsons, T. (1990) *Events in the Semantics of English,* MIT Press.

Teufel, Simone (1999) *Argumentative Zoning: Information Extraction from Scientific Text,* PhD thesis: School of Cognitive Science, University of Edinburgh, Edinburgh, UK.