Translation using Minimal Recursion Semantics

Ann Copestake
University of Stuttgart
Dan Flickinger, Rob Malouf, Susanne Riehemann, Ivan Sag
CSLI, Stanford University
Stanford, CA 94305, USA
{aac,dan,malouf,sr,sag}@csli.stanford.edu

Abstract

We describe minimal recursion semantics (MRS), a framework for semantics within HPSG, which considerably simplifies transfer and generation. We discuss why, in general, a semantic representation with minimal structure is desirable for transfer and illustrate how a descriptively adequate representation with a non-recursive structure may be achieved. The paper illustrates the application of MRS to transfer with a series of examples and compares the approach to others which have been previously adopted within unification based frameworks. Our account involves the use of both language-specific and interlingual predicates or relations and we illustrate how this may be exploited to allow MRS to be used to investigate different lexical semantic approaches.

1 Semantic representation and transfer

In this paper we describe a semantic representation for HPSG known as minimal recursion semantics (MRS), which is being utilized in the English grammar being developed for the Verbmobil project. Verbmobil is a spoken-dialogue machine translation system, which is designed to take German or Japanese input relevant to a limited domain and to produce English output. Our aim is to develop a minimally structured but descriptively adequate representation, which allows for various types of underspecification and facilitates generation and the specification of semantic transfer equivalences. MRS is not, in itself, a full-fledged semantic theory. It can perhaps be best thought of as a meta-level language for describing semantic structures within HPSG. Because MRS supports underspecification, an MRS description will correspond to a set of object language expressions. For simplicity, in the examples in this paper we will take the object language to be predicate calculus, but MRS is intended to be compatible with DRT. The advantages of allowing various types of semantic underspecification for translation purposes are well-known (see e.g. Alshawi et al (1991), Kay et al (1994)), so here we concentrate on the advantages of flatness or minimal structure in a semantic representation language.

The problem of ensuring that a grammar can generate from a particular semantic representation is well-known: it has been discussed in the context of generation by Shieber (1993) and in machine translation by Landsbergen (1987) and Whitelock (1992) among others. One difficulty is the problem of logical form equivalence: even though the grammar may accept a logical form (LF) logically equivalent to a particular LF which is input to the generator, there is no guarantee that it will generate from that syntactic form of the input LF. To take a trivial example, an English grammar might naturally produce the logical form in (1a) from *fierce black cat*, while a straightforward transfer from the natural Spanish representation of *gato negro y feroz* shown in (1b) would produce the LF in (1c), which the English grammar probably would not allow:

(1) a. $\lambda x[\text{fierce}(x) \wedge (\text{black}(x) \wedge \text{cat}(x))]$

b. $\lambda x[\text{gato}(x) \wedge (\text{negro}(x) \wedge \text{feroz}(x))]$

c. $\lambda x[\text{cat}(x) \wedge (\text{black}(x) \wedge \text{fierce}(x))]$

One possible solution to this problem would be for the generator to try all logically equivalent forms. Unfortunately this is not practicable in general since the logical form equivalence problem is undecidable even for first order predicate calculus (Shieber 1993).

There are a variety of possible solutions to this general problem within MT. Attempts have been made to write (or at least tune) the target language grammar so that it accepts the output from transfer. Obviously this approach is inherently non-modular and unidirectional. Landsbergen (1987) proposed that the source and target languages be made isomorphic, with lexical entries and grammar rules put into correspondence in order to guarantee that a derivation in the source language would correspond to one in the target language. This is a principled approach, but requires that task-specific source and target grammars be developed in parallel. As Kay et al (1994) point out, it is quite difficult enough to develop monolingual grammars without imposing the requirement that they be isomorphic to a grammar of another language. This leads to unnatural analyses and necessarily means that the grammars be developed together, specifically for MT purposes. Furthermore, this approach is really only suitable for single language pairs and cannot be used for a project such as Verbmobil, where the requirement is to translate from both German and Japanese into English.

An alternative, which was developed in response to the problems with the parallel grammar approach, is Shake-and-Bake translation (Whitelock 1992, see also Beaven, 1992). As originally described, Shake-and-Bake depends on relating monolingual lexical signs described within a lexicalist grammatical framework. The only information which is actually transferred is the values of the indices which become instantiated during parsing. The 'generator' is given a bag of lexical signs with their semantic arguments instantiated and actually generates by parsing, trying all possible orderings, accepting only those with the appropriate coindexation. The advantage of this approach is that the transfer component contains no information about the monolingual grammar, since it merely relates existing lexical entries, and that the problem of LF equivalence is to a large extent circumvented, because the transfer component only indicates the relation of the lexical signs by coindexation.

One major disadvantage of Shake-and-Bake as originally described is lack of efficiency, since the generator/parser has to consider a number of possibilities which is factorial in the number of signs in the target sentence. Poznanski et al (1995) show that polynomial complexity in Shake-and-Bake generation can be achieved by including extra information from the target language grammar that constrains the possibilities tried by the generator. Although this potentially dilutes the original ideal of independent grammars, it still has advantages over alternative approaches since it could be made clear that this is control information. Even so, Shake-and-Bake (and related systems such as those described by Sanfilippo et al (1992) and Trujillo (1995)) have other disadvantages inherent to strongly lexicalist approaches. Representing phrasal equivalences is cumbersome, since they have to be stated in terms of sets of monolingual entries. The requirement that grammars be absolutely lexicalist restricts the frameworks for which it is appropriate and is even in conflict with recent proposals within HPSG (Sag 1995). As we will illustrate in §3.1, some translation equivalences cannot be encoded simply by relating indices in a conventional semantic framework, and therefore cannot be treated without expansion and complication of the original approach. Finally, translations tend to seem very literal because of the reliance on word to word equivalence and it is difficult to see how a strongly lexicalist approach could accommodate more flexible translations.

It is more usual to insist that transfer produces a LF which is syntactically as well as semantically equivalent to a LF accepted by the generator. Several systems which pursue this approach allow the transfer component to specify both syntax and semantics (e.g. Estival et al, 1990, Arnold and Sadler, 1990) but because we are interested in minimizing the extent to which the writers of the transfer component need knowledge of the target language grammar, we assume that transfer only specifies the semantics. A semantic transfer system can be built without developing special purpose grammars, at least for fairly small systems, as the SRI spoken language translation BCI system (Alshawi, 1992; Alshawi et al 1991) demonstrates. However transfer is much more complex than with Shake-and-Bake and the transfer rules must be written by someone with a detailed knowledge of the target language grammar, since under most proposals for semantic representation the syntactic form of the LF is partially dependent on constituent structure. The BCI system uses semantic transfer at the level of quasi-logical form (QLF): pronouns and scope ambiguities are not resolved, for example. Alshawi et al motivate this decision to use a relatively 'surfacy' representation in detail — the essential point is that ambiguity resolution can be computationally expensive and/or domain dependent and that it should therefore not be carried out unless necessary for translation. However, from our current viewpoint, the QLF representation still retains some unnecessary structure and some examples of transfer have proved difficult to implement in the BCI system.

The approach proposed here can be regarded as taking some elements from a BCI-like approach to semantic transfer and some from Shake-and-Bake (and the variant on that approach described by Trujillo (1995)). Its feasibility rests on two things. The first is the observation that translation is not about strict preservation of denotation. This raises many problems but it does indicate that we do not have to expect that the generator tests all logically equivalent forms of the expression generated by the transfer component, or that the transfer component produces all possible logically equivalent forms acceptable to the target grammar. The second is that, although the theoretical decidability of logical equivalence is independent of the syntax of the representation language, in engineering terms the ease of constructing appropriate forms is highly dependent on the syntax. To return to the example (1) above, the strictly binary nature of $\wedge$ leads to a spurious ambiguity in representation, because the bracketing is irrelevant to the truth conditions. But the generator can only determine this by examining the logical properties of $\wedge$. An alternative representation could have been adopted which does not introduce spurious semantic ambiguity by redundant bracketing. For example, instead of (1a) (repeated as (2a)) we could have written (2b) where $\wedge$ takes a set of arguments. To avoid the computational expense of deciding whether two sets have the same members we could agree to give this the canonical form shown in (2c) where a conventional order is imposed on the propositions.

(2) a. $\lambda x[\text{fierce}(x) \wedge (\text{black}(x) \wedge \text{cat}(x))]$

b. $\lambda x[\bigwedge\{\text{fierce}(x), \text{black}(x), \text{cat}(x)\}]$

c. $\lambda x[\bigwedge\{\text{black}(x), \text{cat}(x), \text{fierce}(x)\}]$

Consider the effect this has on transfer from Spanish to English. Instead of a series of recursively applied rules which either overgenerate with concomitant lack of efficiency or have to be constructed with knowledge of the English grammar of adjectives, we can have a simple, non recursive procedure which generates (2c) purely on the basis of the simple transfer equivalences $\text{black}(x) \leftrightarrow \text{negro}(x)$, $\text{fierce}(x) \leftrightarrow \text{feroz}(x)$ and $\text{cat}(x) \leftrightarrow \text{gato}(x)$. More importantly, it becomes much easier to state those equivalences for which there is no one to one mapping between lexically realized predicates. For example, for the equivalence English:*young bull* $\approx$ Spanish:*novillo*, we can simply state $\{\text{bull}(x), \text{young}(x)\} \leftrightarrow \{\text{novillo}(x)\}$, without having to worry about the possibility of interpolated material (e.g. *young black bull*).

In general, the greater the structural complexity of the representation, the worse the problem of constructing a LF acceptable to the target grammar becomes. A separate, though not unrelated, issue is that the semantic representation chosen should allow for underspecification of the sorts of information which are not straightforwardly resolvable when analyzing natural languages, where the ambiguity can be naturally preserved on translation.[1] The problem then is to define a semantic representation which is logically adequate but not unnecessarily structured, which reflects the properties that translation should aim to preserve and which allows for underspecification. A final essential criterion is that the representation should be capable of supporting inference.

The proposal for the translation mechanism made here can be classified as semantic transfer.[2] It relies on the transfer component constructing a representation which can be accepted by the target grammar, with some limited contribution from the generator. The transfer component will sometimes output multiple forms, some of which may be unacceptable to the target grammar (a technique referred to as filtering in the BCI system). There will frequently be multiple possible inputs to the generator, but these will be ordered by a control mechanism which is distinct from the declarative transfer equivalence specification. Further details are given in the rest of this paper, after the description of MRS.

[1] For English/German, this is often true for quantifier scope and PP attachment ambiguities, but does not, in general, apply to lexical ambiguity, except for some semi-regular cases of polysemy which hold cross-linguistically, such as the abstract representation/physical object sense of *newspaper*.

[2] This is a matter of terminology, since there is no sharp dividing line between semantic transfer and those interlingual approaches which do not insist on the identity of source and target language representations. We will assume that some predicates are interlingual and some language-specific in what follows.

2 An introduction to MRS

We begin by discussing some general issues in achieving a minimally structured semantic representation. We will assume a neo-Davidsonian style of representation with explicit event variables. This immediately leads to a rather flat representation style: contrast (3b,c) with (3d) (these examples and others in this introductory section are for expository purposes, not a serious proposal for the representation of temporal adverbials etc):

(3) a. On Monday Kim ran in Foothills Park.

b. on(in(run(Kim), Foothills_Park), Monday)

c. in(on(run(Kim), Monday), Foothills_Park)

d. on(e, Monday) $\wedge$ run(e, Kim) $\wedge$ in(e, Foothills_Park)

The structure in (3d) is more appropriate because it does not introduce the spurious scoping distinction.

The most important use of recursive structures in a semantic representation is to allow for scope. For example, consider the representation of:

(4) Every tall man is old.

In this case there is only one possible scope for *every*, which is shown in (5) using generalized quantifiers:

(5) $\mathrm{every}(x, \mathrm{man}(x) \wedge \mathrm{tall}(x), \mathrm{old}(x))$

We should, therefore, be able to retrieve this reading unambiguously from the semantic representation that the grammar constructs for this sentence. However, if we have the totally flat structure shown in (6) it is impossible to retrieve the correct reading unambiguously, because we would get the same structure for (7).

(6) $\bigwedge\{\mathrm{every}(x), \mathrm{man}(x), \mathrm{old}(x), \mathrm{tall}(x)\}$

(7) Every old man is tall.

We therefore require a flat representation which preserves sufficient information about the scope of a quantifier to be able to construct all and only the possible readings for the sentence. Similar remarks apply with respect to the representation of the scope of *not*, *or* and so on.

We can achieve this effect while retaining the flatness of the representation by adding extra variables, which have the effect of capturing scope information. For example, we could represent (7) as (8):

(8) $\bigwedge\{\mathrm{every}(x, 1, 2), \mathrm{man}_1(x), \mathrm{old}_1(x), \mathrm{tall}_2(x)\}$

These extra variables can be thought of as handles which enable us to 'grab' particular propositions in the flat list (cf the use of *labels* in Frank and Reyle, 1994). In (8) the scoped representation can be reconstructed by replacing the handles in the restriction and body arguments of *every* with the propositions tagged by those handles. From now on we will drop the use of $\wedge$ and assume implicit conjunction.

In the example above, we simply tagged everything which would have been inside a set of braces in a conventional formula with the same subscript. Nested quantifiers do not require multiple indices on a single conjunct, since we can trace the nesting via the restriction of the embedded quantifier. The reason why this approach is not just a notational variant of a standard representation is that the handles can be underspecified to represent multiple scopes. For example, the underspecified representation of (9a) would be (9b):[3]

(9) a. Every dog chased some cat.

b. $\text{every}_1(x, 3, n), \text{dog}_3(x), \text{cat}_7(y), \text{some}_5(y, 7, m), \text{chase}_4(e, x, y)$

Here n and m stand for variables over handles. The scoped representations would be:

(10) a. $\text{every}_1(x, 3, 4), \text{dog}_3(x), \text{cat}_7(y), \text{some}_5(y, 7, 1), \text{chase}_4(e, x, y)$
(wide scope *some*)

b. $\text{every}_1(x, 3, 5), \text{dog}_3(x), \text{cat}_7(y), \text{some}_5(y, 7, 4), \text{chase}_4(e, x, y)$
(wide scope *every*)

MRS proper is defined in terms of feature structures (FSs), rather than the linearized representation shown above. The semantic representation has two parts, CONTENT and CONTEXT, as usual in HPSG, but here we are mainly concerned with the CONTENT value. An MRS expression consists of a structure of sort *mrs-struct*, with appropriate features HANDEL and LISZT, which take values of sort *handle* and *list* respectively. We use the feature name LISZT to distinguish the non-recursive semantic structure from ordinary lists: the values of LISZTs have to be treated like sets in some respects as we will see below and LISZT is an appropriate name for an extraordinary object which composes the semantics. In keeping with this theme, since handle sounds rather plebeian, we use HANDEL for the other main feature of the compositional semantics. We adopt the normal convention of writing feature names in (small) capitals and sorts/types in italics. In what follows, we will often use liszt and handel in the normal lower case font to refer to the values of those features.

The sort *mrs-struct* also has the appropriate feature INDEX which plays much the same role as a lambda variable in conventional representations. The value of LISZT is defined to be a flat list of *rels* (relations), which all have HANDELs and other features depending on their sort. As in Pollard and Sag (1994:chap. 8, sec. 5), the actual relation is indicated by the sort of the rel. Determiners, such as *every*, have rels with appropriate features BV (bound variable) which takes a value of sort *ref-ind*, and RESTR (restriction) which takes a *handle*. Verb rels have a feature EVENT which takes an event variable. They have features such as ACT (actor) and UND (undergoer) following Davis (forthcoming). Common nouns have rels with the feature INST which takes a *ref-ind*. We will introduce some other specialized subsorts of rel later in this document. An example of the (unscoped) MRS representation for (9a) is shown in Figure 1. Here the variable sorts and the internal structure of the indices are not shown, but only the coindexation between them,

[3]We are oversimplifying somewhat here, since in order to be able to represent any information about the relative scope of more than two quantifiers, we need to be able to represent the relative scope of pairs of quantifiers (Frank and Reyle, 1994). This can be accomodated in the MRS representation, but we do not give the details since there are not relevant here.

$$\begin{bmatrix} \text{HANDEL} & \boxed{1} \vee \boxed{5} \\ \text{INDEX} & \boxed{8} \\ \text{LISZT} & \left\langle {}_{every\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{1} \\ \text{BV} & \boxed{2} \\ \text{RESTRICTION} & \boxed{3} \end{bmatrix}, {}_{dog\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{3} \\ \text{INST} & \boxed{2} \end{bmatrix}, {}_{some\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{5} \\ \text{BV} & \boxed{6} \\ \text{RESTRICTION} & \boxed{7} \end{bmatrix}, \right. \\ & \left. {}_{cat\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{7} \\ \text{INST} & \boxed{6} \end{bmatrix}, {}_{chase\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{4} \\ \text{EVENT} & \boxed{8} \\ \text{ACT} & \boxed{2} \\ \text{UND} & \boxed{6} \end{bmatrix} \right\rangle \end{bmatrix}$$

Figure 1: Unscoped representation for *every dog chased some cat*

$$\begin{bmatrix} \text{HANDEL} & \boxed{5} \\ \text{INDEX} & \boxed{8} \\ \text{LISZT} & \left\langle {}_{every\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{1} \\ \text{BV} & \boxed{2} \\ \text{RESTRICTION} & \boxed{3} \\ \text{BODY} & \boxed{4} \end{bmatrix}, {}_{dog\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{3} \\ \text{INST} & \boxed{2} \end{bmatrix}, {}_{some\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{5} \\ \text{BV} & \boxed{6} \\ \text{RESTRICTION} & \boxed{7} \\ \text{BODY} & \boxed{1} \end{bmatrix}, \right. \\ & \left. {}_{cat\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{7} \\ \text{INST} & \boxed{6} \end{bmatrix}, {}_{chase\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{4} \\ \text{EVENT} & \boxed{8} \\ \text{ACT} & \boxed{2} \\ \text{UND} & \boxed{6} \end{bmatrix} \right\rangle \end{bmatrix}$$

Figure 2: MRS for *every dog chased some cat* with *some* taking wide scope

indicated as usual by boxed integers. The handel shown at the outer level allows the sentence to be embedded, as in *Sandy said that every dog chased some cat*, for example. Here it is a disjunct of the handels of the quantifiers, because we have not assigned a scope: in the scoped representation the handel will be the handel of the widest scoped quantifier. Scoped quantifiers have to be represented as having both a restriction and a body. However, for the underspecified representation the body of the quantifiers is left unspecified. The representation can be monotonically enriched to either scoped structure by appropriate coindexation of the handels and instantiation of the BODYs of the quantifiers. One of the scoped representations (wide scope *some*) is shown in Figure 2. Further details are given in Copestake et al (1995).

From now on, rather than simply using boxed integers to indicate reentrancy in the conventional way, we will indicate their types for clarity, using h for handel, e for event, x for entity, and i for individual (event or entity). The MRS representation of (11) is shown

$$
\begin{bmatrix}
\text{HANDEL} & \boxed{h1} \\
\text{INDEX} & \boxed{e8} \\
\text{LISZT} & \left\langle {}_{good\_rel}\begin{bmatrix}\text{HANDEL} & \boxed{h4} \\ \text{ARG} & \boxed{e8}\end{bmatrix}, {}_{neg\_rel}\begin{bmatrix}\text{HANDEL} & \boxed{h3} \\ \text{INST} & \boxed{h4}\end{bmatrix}, {}_{suit\_rel}\begin{bmatrix}\text{HANDEL} & \boxed{h5} \\ \text{EVENT} & \boxed{e8} \\ \text{ACT} & \boxed{i7} \\ \text{UND} & \boxed{x6}\end{bmatrix}, \right. \\
 & \left. {}_{ego\_rel}\begin{bmatrix}\text{HANDEL} & \boxed{h9} \\ \text{INST} & \boxed{x6}\end{bmatrix}, {}_{deixis\_rel}\begin{bmatrix}\text{HANDEL} & \boxed{h10} \\ \text{INST} & \boxed{i7}\end{bmatrix}, {}_{real\_rel}\begin{bmatrix}\text{HANDEL} & \boxed{h2} \\ \text{ARG} & \boxed{h3}\end{bmatrix} \right\rangle
\end{bmatrix}
$$

Figure 3: MRS for *That really doesn't suit me well*

in Figure 3: we will use this example in the discussion of translation in the next section.

(11) That really doesn't suit me well

Here the handels of rels which are not scoped with respect to one another are unified: in general whenever two *mrs-struct*s are combined, we unify the handels to create the handel for the result unless a rel in the liszt of one *mrs-struct* takes the handel of the other as an argument. Becaause scope has not been assigned, the outermost handel is not coindexed with any of the handels internal to the rels — it should be thought of as a disjunction of h2 and h10. The new liszt is constructed by appending the liszts of the structures.

The inclusion of the unresolved *ego_rel* and *deixis_rel* reflects the idea that the semantics should preserve enough surface information to facilitate translation for cases where there is a straightforward mapping between languages. This is not built into MRS in any way, however, so the representation is more flexible than QLF or Shake-and-Bake. Note that *well* has *good_rel* in its semantics and that *really* contains a *real_rel*. For adverbs where there is a systematic distinction in meaning from the corresponding adjective it would be necessary to use a separate rel. The current treatment does not preclude *suit*, *real* and so on from being polysemous in that there could be subsorts of *suit_rel*, *real_rel* and so on. The nominal senses of *suit* will have different rels, of course. The event e8 is actually a state, but we will not draw any distinctions between eventuality types for the purposes of this paper. We also omit any representation of tense here, since MRS is compatible with a variety of approaches. Note that both *real_rel* and *neg_rel* take handels as arguments to allow for the difference between (12a) and (12b):

(12) a. That really doesn't suit me

b. That doesn't really suit me

This contrasts with PPs such as *on Monday* and *in Foothill Park* which are represented as taking events to avoid spurious ambiguity. Similarly we assume that tense rels take event arguments.

3 Translation

Translation between MRS representations basically depends on setting up the correct coindexations between the arguments of the rels.[4] A trivial example is shown in the Figure 4, where all the mappings between relations are one-to-one. In this section, we will use the convention that German MRS expressions are shown on the left of the transfer equivalence and English on the right. The equivalences are all symmetric and bidirectional, however. By making transfer equivalences themselves be FSs, we can use reentrancies between the two halves of the equivalences to indicate argument identification and specify generalizations about classes of transfer equivalence using sorts. Note that the output MRS is underspecified, in that the top-level HANDEL and INDEX values are not instantiated, but constraints on the English grammar would mean that h1 and e4 are the only possible values.

Most semantic transfer systems require *structural* transfer rules which handle the reordering and rearrangement of the semantic structures. Structural transfer rules are minimized by the MRS representation, which is desirable not only because of the reduction in number and complexity of the transfer rules, but because it is extremely difficult to guarantee that phrasal transfer rules and structural rules will interact properly. Transfer between two MRS representations essentially requires a single structural transfer rule, if we think of the value of LISZT as a set of rels:

Definition 1 (Structural transfer rule) *The translation τ of a set of rels $z = x \cup y$ is defined as*

$$\tau(x \cup y) = \tau_l(x) \cup \tau(y)$$

where τ_l is the base case: x translates as $\tau_l(x)$ iff there is some transfer equivalence such that the input unifies with x giving $\tau_l(x)$ as the output.

In other words, the translation of a *mrs-struct* is the translation of some subset of the liszt unioned with the translation of the remainder of the liszt.[5]

3.1 Transfer equivalences

All transfer equivalences consist of relationships between sets of rels but the vast majority will be lexically motivated. We can distinguish two subcases: lexical transfer equivalences, where the liszts which are related by the equivalence will each be found in a lexical sign, and phrasal transfer equivalences, where one or more of the liszts can only be realized as (part of) a phrasal sign. For example, the following are examples of lexical transfer equivalences that could apply in the translation of (13a) (which has the MRS representation shown in Figure 5) to (12) (repeated as (13b)):

[4] For simplicity, we show coindexation between the complete indices here, but the true representation is slightly more complex, since in HPSG indices contain information which will not be shared between languages, such as gender values, for example.

[5] Although this rule could be regarded simply as a constraint between input and output structures, it has to be implemented specially within constraint based formalisms which lack a representation of sets.

Input: *Kim singt*

Input MRS:

$$\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{INDEX} & \boxed{e4} \\ \text{LISZT} & \left\langle {}_{named\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{NAME} & Kim \\ \text{NAMED} & \boxed{x2} \end{bmatrix}, {}_{singen\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{EVENT} & \boxed{e4} \\ \text{ACT} & \boxed{x2} \end{bmatrix} \right\rangle \end{bmatrix}$$

Transfer equivalences:

$${}_{singen\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h2} \\ \text{EVENT} & \boxed{e3} \\ \text{ACT} & \boxed{x1} \end{bmatrix} \Longleftrightarrow {}_{sing\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h2} \\ \text{EVENT} & \boxed{e3} \\ \text{ACT} & \boxed{x1} \end{bmatrix}$$

$${}_{named\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h3} \\ \text{NAME} & \boxed{n1} \\ \text{NAMED} & \boxed{x4} \end{bmatrix} \Longleftrightarrow {}_{named\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h3} \\ \text{NAME} & \boxed{n1} \\ \text{NAMED} & \boxed{x4} \end{bmatrix}$$

Instantiated equivalences:

$${}_{singen\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{EVENT} & \boxed{e4} \\ \text{ACT} & \boxed{x2} \end{bmatrix} \Longleftrightarrow {}_{sing\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{EVENT} & \boxed{e4} \\ \text{ACT} & \boxed{x2} \end{bmatrix}$$

$${}_{named\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{NAME} & \boxed{n1}\ Kim \\ \text{NAMED} & \boxed{x2} \end{bmatrix} \Longleftrightarrow {}_{named\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{NAME} & \boxed{n1} \\ \text{NAMED} & \boxed{x2} \end{bmatrix}$$

Output MRS:

$$\begin{bmatrix} \text{HANDEL} & \\ \text{INDEX} & \\ \text{LISZT} & \left\langle {}_{named\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{NAME} & Kim \\ \text{NAMED} & \boxed{x2} \end{bmatrix}, {}_{sing\_rel}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{EVENT} & \boxed{e4} \\ \text{ACT} & \boxed{x2} \end{bmatrix} \right\rangle \end{bmatrix}$$

Output: *Kim sings*

Figure 4: Transfer using MRS

$$\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{INDEX} & \boxed{e8} \\ \text{LISZT} & \left\langle {}_{\mathit{schlecht\_rel}}\begin{bmatrix}\text{HANDEL} & \boxed{h3}\\ \text{ARG} & \boxed{e8}\end{bmatrix}, {}_{\mathit{passen\_rel}}\begin{bmatrix}\text{HANDEL} & \boxed{h5}\\ \text{EVENT} & \boxed{e8}\\ \text{THEME} & \boxed{i7}\\ \text{EXP} & \boxed{x6}\end{bmatrix}, \right. \\ & \left. {}_{\mathit{ego\_rel}}\begin{bmatrix}\text{HANDEL} & \boxed{h9}\\ \text{INST} & \boxed{x6}\end{bmatrix}, {}_{\mathit{deixis\_rel}}\begin{bmatrix}\text{HANDEL} & \boxed{h10}\\ \text{INST} & \boxed{i7}\end{bmatrix}, {}_{\mathit{echt\_rel}}\begin{bmatrix}\text{HANDEL} & \boxed{h2}\\ \text{ARG} & \boxed{h3}\end{bmatrix} \right\rangle \end{bmatrix}$$

Figure 5: MRS for *Das paßt echt schlecht bei mir*

(13) a. Das paßt echt schlecht bei mir

b. That really doesn't suit me well

$$\left\langle {}_{\mathit{passen\_rel}}\begin{bmatrix}\text{THEME} & \boxed{i1}\\ \text{EXP} & \boxed{x2}\end{bmatrix} \right\rangle \Longleftrightarrow \left\langle {}_{\mathit{suit\_rel}}\begin{bmatrix}\text{ACT} & \boxed{i1}\\ \text{UND} & \boxed{x2}\end{bmatrix} \right\rangle$$

$$\left\langle {}_{\mathit{echt\_rel}}\begin{bmatrix}\text{ARG} & \boxed{h1}\end{bmatrix} \right\rangle \Longleftrightarrow \left\langle {}_{\mathit{real\_rel}}\begin{bmatrix}\text{ARG} & \boxed{h1}\end{bmatrix} \right\rangle$$

Note that we have refrained from relating the non-argument handels and event variables explicitly here. It is redundant to repeat the information that the handels and events are in one-to-one relationship for each rel, since it could simply be inherited from a general sort for transfer equivalences (in a manner comparable to the use of *tlinks* described by Copestake and Sanfilippo (1993) and Sanfilippo et al (1992)). But in certain examples, event variables may not correspond one-to-one across languages. We will discuss one case in §3.3.

A simple phrasal transfer equivalence, for *schlecht*≈*not good* is given below:[6]

$$\left\langle {}_{\mathit{schlecht\_rel}}\begin{bmatrix}\text{ARG} & \boxed{e1}\ \mathit{event}\end{bmatrix} \right\rangle \Longleftrightarrow \left\langle {}_{\mathit{neg\_rel}}\begin{bmatrix}\text{ARG} & \boxed{h2}\end{bmatrix}, {}_{\mathit{good\_rel}}\begin{bmatrix}\text{HANDEL} & \boxed{h2}\\ \text{ARG} & \boxed{e1}\end{bmatrix} \right\rangle$$

The rule would apply to both adjectival and adverbial uses of *schlecht* (we assume that *good* and *well* both have *good_rel* in their semantics). However, we have specialized the antecedent so that it only applies when an event is modified.

Note that the use of handels allows a transfer relationship to be expressed by relating variables in a way which would be impossible with a more conventional semantic representation without adopting additional devices (such as transfer variables). Shake-and-Bake

[6]There will also be an equivalence *schlecht*≈*bad*, of course, but the equivalence with *not good* seems justified by the regularity with which *schlecht* is used in contexts where *bad* is an inappropriate translation. We discuss how we avoid the translation of (12a) as **that really suits me badly* in §3.2.

as described in Whitelock (1992) and Beaven (1992) could not represent this example adequately because it does not have transfer variables and the semantic representation assumed has nothing equivalent to handels.

Some rels such as *neg_rel*, *ego_rel* and *deixis_rel* are not language-specific — in general we allow relations with a sort which is defined to be interlingual to simply be transferred as they stand between source and target.

3.2 Generation and control

Since transfer produces a logical form, standard techniques for generating from a semantic representation can be used with MRS. Lexical entries can be indexed by rel name for efficient lookup. If the grammar and transfer were completely lexicalist — i.e. if the rels generated by transfer corresponded directly to single lexical entries, then the translation process could actually be implemented in a manner very similar to Shake-and-Bake, since unifying the rels with the semantics of the relevant lexical entry would give an instantiated lexical sign. However, because MRS decouples the lexicon and the semantics, more flexible strategies are possible, while retaining the advantages of lexicalism where it is appropriate. The creation of a complete target logical form prior to generation of the target string has the advantage over the Shake-and-Bake style approach that a strategic generator can check the logical form and either request an alternative from transfer or rewrite it itself, if that appears appropriate.

For instance, taking the example above, we wish to avoid the generation of (14):

(14) * That really suits me badly

The semantic representation corresponding to this will be output by the transfer equivalences, but it is excluded on monolingual grounds, since verbs which denote a positive attitude towards something such as *suit*, *like* and so on cannot be modified by negative adverbials (Condoravdi and Sanfilippo, 1990). Since for this case the unacceptability can be determined by examination of the target logical form alone, the filtering could be done prior to generation proper for efficiency.

The transfer equivalences given above allow *schlecht* to be translated as *not well* in any context where it is modifying an event. This is sometimes desirable, but usually should be blocked by the translation as *badly*. As a general heuristic, we can specify that translations which produce single lexical items are to be preferred over those which produce phrases. This selection could reasonably be part of the functionality of the generator, rather than the transfer component, since it essentially corresponds to the monolingual principle of avoiding periphrasis and is dependent on the number of lexemes involved, not on the number of rels. But, since in general we assume that the transfer component incorporates heuristics which order translation equivalences, it might be more practical to include a heuristic which avoids equivalences which have phrases on the target side if alternative lexical outputs are available.

3.3 A more complex example

To facilitate comparison with other approaches to transfer, we now consider the standard example of the translation of English verbs of manner of motion (e.g. *swim*, *stagger*) occurring with a locative expression describing a completed path. In many other languages, including the Romance languages and Japanese, this pattern is not possible (Talmy 1985). Here we will concentrate on Spanish, where the manner of motion is expressed with an adverbial, while the main verb conveys the path. For example, (15a) is usually translated as (15b) (which is more literally translated as (15c)).

(15) a. Kim swam across the river
  b. Kim cruzó el río nadando
  c. Kim crossed the river swimming

In order to give a direct comparison between transfer using MRS and an alternative sign-based approach, we will use a comparable semantic analysis to that given by Sanfilippo et al (1992).[7] In this analysis, the gerundive is translated as expressing a distinct event, which is related to the main verb event by a predicate **while**. For current purposes, the distinction between MRS and Sanfilippo et al's encoding of a neo-Davidsonian semantics by use of proto-roles is not significant. We show a linearized equivalent of their encoding of the semantics for *nadando*[8] and the corresponding MRS representation below:

$$[e2][\text{while}(e2, e1) \wedge \text{nadar}(e2) \wedge \text{p-agt-cause-move-manner}(e2, x1) \wedge P(e1)]$$

$$\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{INDEX} & \boxed{e2} \\ \text{LISZT} & \left\langle {}_{\mathit{while\_rel}}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{ARG1} & \boxed{e2} \\ \text{ARG2} & \boxed{e1} \end{bmatrix}, {}_{\mathit{nadar\_rel}}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{EVENT} & \boxed{e2} \\ \text{ACT} & \boxed{x1} \end{bmatrix}, {}_{\mathit{rel}}\begin{bmatrix} \text{HANDEL} & \boxed{h1} \\ \text{EVENT} & \boxed{e1} \end{bmatrix} \right\rangle \end{bmatrix}$$

A simplified form of the MRS representations for (15a) and (15b) are shown in Figure 6. We assume that strict intransitive manner of motion verbs such as *swim* are related to entries subcategorized for a PP expressing a path by lexical rule. The details of this are unimportant here, but we assume that in Spanish this rule cannot apply to manner of motion verbs when the path is bounded. In contrast to Sanfilippo et al's use of a Shake-and-Bake style representation of transfer equivalence between sets of lexical signs, we make use of the transfer equivalences between liszts shown below (in this section, we are using the convention that the English rels are on the left of the translation equivalence).

[7]For some other approaches to this mismatch problem see, for example, Isabelle et al, 1988; Beaven, 1992; Nirenburg and Levin, 1992; Dorr, 1994.

[8]We have omitted the classification of events which Sanfilippo et al adopt — this is an important part of the analysis but is not relevant to the comparison with MRS since we could implement it in exactly the same way, by typing event variables. Apart from this, we have essentially preserved the analysis, but we do not reproduce the original FS here, since it is multiply nested, due to the use of a binary **and** and over 20 lines of text.

$$
\left[\begin{array}{ll}
\text{HANDEL} & \boxed{h1} \\
\text{INDEX} & \boxed{e1} \\
\text{LISZT} & \left\langle {}_{named\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h1} \\ \text{NAME} & Kim \\ \text{NAMED} & \boxed{x1}\end{array}\right], {}_{swim\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h1} \\ \text{EVENT} & \boxed{e1} \\ \text{ACT} & \boxed{x1}\end{array}\right], \right. \\
 & \left. {}_{across\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h1} \\ \text{ARG1} & \boxed{e1} \\ \text{ARG2} & \boxed{x2}\end{array}\right], {}_{the\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h1} \\ \text{BV} & \boxed{x2} \\ \text{RESTR} & \boxed{h2}\end{array}\right], {}_{river\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h2} \\ \text{INST} & \boxed{x2}\end{array}\right] \right\rangle
\end{array}\right]
$$

$$
\left[\begin{array}{ll}
\text{HANDEL} & \boxed{h1} \\
\text{INDEX} & \boxed{e1} \\
\text{LISZT} & \left\langle {}_{named\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h1} \\ \text{NAME} & Kim \\ \text{NAMED} & \boxed{x1}\end{array}\right], {}_{while\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h1} \\ \text{ARG1} & \boxed{e2} \\ \text{ARG2} & \boxed{e1}\end{array}\right], {}_{nadar\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h1} \\ \text{EVENT} & \boxed{e2} \\ \text{ACT} & \boxed{x1}\end{array}\right], \right. \\
 & \left. {}_{cruzar\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h1} \\ \text{EVENT} & \boxed{e1} \\ \text{ACT} & \boxed{x1} \\ \text{PATH} & \boxed{x2}\end{array}\right], {}_{el\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h1} \\ \text{BV} & \boxed{x2} \\ \text{RESTR} & \boxed{h2}\end{array}\right], {}_{rio\_rel}\!\left[\begin{array}{ll}\text{HANDEL} & \boxed{h2} \\ \text{INST} & \boxed{x2}\end{array}\right] \right\rangle
\end{array}\right]
$$

Figure 6: Simplified MRS representations for *Kim swam across the river* and *Kim cruzó el río nadando*

$$\left\langle {}_{swim\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix}, {}_{path\_rel}\begin{bmatrix}\text{ARG1} & \boxed{e1}\end{bmatrix}\right\rangle \Longleftrightarrow$$

$$\left\langle {}_{nadar\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e2}\\ \text{ACT} & \boxed{x1}\end{bmatrix}, {}_{while\_rel}\begin{bmatrix}\text{ARG1} & \boxed{e2}\\ \text{ARG2} & \boxed{e1}\end{bmatrix}, {}_{move\_path\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix}\right\rangle$$

$$\left\langle {}_{across\_rel}\begin{bmatrix}\text{ARG1} & \boxed{e1}\\ \text{ARG2} & \boxed{x2}\end{bmatrix}, {}_{move\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix}\right\rangle \Longleftrightarrow \left\langle {}_{cruzar\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\\ \text{PATH} & \boxed{x2}\end{bmatrix}\right\rangle$$

In these rules we have made use of the underspecified rels *path_rel*, *move_path_rel* and *move_rel* which subsume *across_rel*, *cruzar_rel* and *swim_rel* respectively. These constrain the transfer equivalences in a general manner, so that the former equivalence will apply to a range of paths and the latter to a range of manner of movement verbs. The easiest way to understand these underspecified equivalences is to assume that they are unified in advance of transfer, to give a fully instantiated transfer equivalence, as shown below:

$$\left\langle {}_{swim\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix}, {}_{across\_rel}\begin{bmatrix}\text{ARG1} & \boxed{e1}\\ \text{ARG2} & \boxed{x2}\end{bmatrix}\right\rangle \Longleftrightarrow$$

$$\left\langle {}_{nadar\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e2}\\ \text{ACT} & \boxed{x1}\end{bmatrix}, {}_{while\_rel}\begin{bmatrix}\text{ARG1} & \boxed{e2}\\ \text{ARG2} & \boxed{e1}\end{bmatrix}, {}_{cruzar\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\\ \text{PATH} & \boxed{x2}\end{bmatrix}\right\rangle$$

Together with the straightforward equivalences between *river* and *rio* etc, this enables (15a) to be translated as (15b) (or vice versa).

The transfer equivalence for *swim* can be generated automatically from the usual equivalence between *swim_rel* and *nadar_rel* shown below.

$$\left\langle {}_{swim\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix}\right\rangle \Longleftrightarrow \left\langle {}_{nadar\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix}\right\rangle$$

This is achieved by making use of a metalevel pattern (shown below) along the lines of the *tlink-rules* described by Copestake and Sanfilippo (1993), which will also produce the more complex equivalence from other manner of movement verbs equivalences such as English:*float*≈Spanish:*flotar*. Here *manner_move_rel* is a sort which subsumes *swim_rel*, *float_rel*, *nadar_rel* etc and thus the translation equivalence between *swim_rel* and *nadar_rel* will instantiate the top half of the structure, generating the new equivalence.

$$\left\langle {}_{manner\_move\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix}\right\rangle \Longleftrightarrow \left\langle {}_{manner\_move\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix}\right\rangle$$

$$\Downarrow$$

$$\left\langle {}_{manner\_move\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix}, {}_{path\_rel}\begin{bmatrix}\text{ARG1} & \boxed{e1}\end{bmatrix}\right\rangle \Longleftrightarrow$$

$$\left\langle {}_{manner\_move\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e2}\\ \text{ACT} & \boxed{x1}\end{bmatrix}, {}_{move\_path\_rel}\begin{bmatrix}\text{EVENT} & \boxed{e1}\\ \text{ACT} & \boxed{x1}\end{bmatrix} {}_{while\_rel}\begin{bmatrix}\text{ARG1} & \boxed{e2}\\ \text{ARG2} & \boxed{e1}\end{bmatrix}\right\rangle$$

It is essentially this meta-pattern which encapsulates the generalization about translation. Like a tlink-rule, it can be implemented straightforwardly within a unification based framework if transfer equivalences are themselves FSs.

The MRS representation is considerably simpler than that of Sanfilippo et al, though for reasons of space we do not show their transfer equivalences here. Unlike that approach, it does not specify the grammatical means employed in Spanish to create the adverbial manner component. Alternative semantic analyses are possible, however. It might be preferable to introduce an interlingual notion of paths, since there seem to be a limited number of possibilities for lexically encoded paths, which would allow us to avoid expressing the transfer equivalence between *cruzar* and *across*. Furthermore, it would be possible to analyze English sentences such as (15a) as having two event variables. This would give a semantics more similar to that proposed by Talmy (1985). Thus we could take an approach within MRS which was closer to an interlingual representation (see e.g. Nirenburg and Levin (1992) and Dorr (1994) for a more interlingual approach to this example), if monolingual criteria seemed to justify the move. We will not go into details here but we believe that our use of MRS makes it much more straightforward to concentrate on the purely semantic aspects of representation and transfer problems and hence to investigate a variety of approaches to lexical semantics.

4 Conclusion

We have shown here how MRS representations simplify the description of transfer equivalences and alleviate the problem of ensuring that the grammar generates from the output of transfer. The potential cost of doing this is in efficiency, since the generator cannot take advantage of any clues about the syntax that may be implicit in a more structured representation. However, we believe that control statements should be stated explicitly, both to avoid complicating the transfer component and because this will lead to greater efficiency for a particular generation strategy. There will normally be several possible sentences that can be generated from any one semantic representation output by the transfer component. In principle this is desirable to avoid stilted translations: sentences would be ordered by considering monolingual information such as preferred collocations and constructions and by explicitly identifying the discourse effects that affect word order. In practice we would expect that heuristics would be used (such as preferring the word order that corresponds most closely with the source language).

As we mentioned in §3.2, the approach to transfer described here would be similar to Shake-and-Bake, if there were a one-to-one relationship between rels and lexical items. But as we have illustrated here, the MRS approach allows a less cumbersome treatment of some phrasal translation equivalences and facilitates investigating representations which are partially interlingual. We also believe that there will be long-term advantages in avoiding direct linkage between semantics and lexical items, because it makes it possible to incorporate more flexible translation and generation strategies, for the cases where a lexical approach gives stilted or unnatural results. This is currently most relevant for domain-specific speech-based MT systems, such as Verbmobil, where it is feasible to incorporate at least limited inference on a domain model, and where it is especially important to be able to generate natural-sounding output.

As we said in the introduction, MRS is intended as a metalevel representation which is compatible with semantic theories such as DRT. In this paper, we have concentrated on illustrating the advantages of a flat semantics for transfer, but MRS also allows underspecification of quantifier scope and PP attachment (by underspecification of arguments) and of lexical ambiguity (by use of a hierarchy of rel types).

Acknowledgements

We are grateful to Martin Kay, Christian Rohrer and Arturo Trujillo for ideas and comments on this paper. This work was partially funded by the German Federal Ministry for Research and Technology (BMFT) in the framework of the Verbmobil Project. The responsibility for the contents of this study lies with the authors.

References

Alshawi, H. (ed.) (1992) *The Core Language Engine,* MIT Press, Cambridge, Mass.

Alshawi, H., D. Carter, M. Rayner and B. Gambäck (1991) 'Translation by Quasi Logical Form transfer', *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics (ACL-91),* Berkeley, California, pp. 161–168.

Arnold, D. and L. Sadler (1990) 'The theoretical basis of MiMo', *Machine Translation,* **5 (3)**, 195–222.

Beaven, J. L. (1992) 'Lexicalist unification-based machine translation', PhD thesis, University of Edinburgh.

Condoravdi, C. and A. Sanfilippo (1990) 'Notes on Psychological Predicates', ESPRIT BRA-3030 ACQUILEX working paper 7: University of Cambridge Computer Laboratory.

Copestake, A. and A. Sanfilippo (1993) 'Multilingual lexical representation', *Proceedings of the AAAI Spring Symposium on Building lexicons for machine translation,* Stanford, CA.

Copestake, A., D. Flickinger, R. Malouf, S. Riehemann, I. Sag (1995) 'Minimal Recursion Semantics', ms. CSLI, Stanford University.

Davis, A. (forthcoming) 'Linking and the hierarchical lexicon', Doctoral dissertation, Stanford University.

Dorr, B.J. (1994) 'Machine Translation Divergences: a formal description and a proposed solution', *Computational Linguistics,* **20 (4)**, 597–633.

Estival, D., A. Ballim, G. Russell and S. Warwick (1990) 'A syntax and semantics for feature structure transfer', *Proceedings of the 3rd International Conference on theoretical and methodological issues in MT of NLs (TMI-90),* Austin,Texas, pp. 131–143.

Frank, A. and U. Reyle (1994) 'Principle based semantics for HPSG', to appear in: Arbeitspapiere des Sonderforschungsbereichs 340, University of Stuttgart.

Isabelle, P., M. Dymetman and E. Macklovitch (1988) 'CRITTER: a translation system for agricultural market reports', *Proceedings of the 12th International Conference on Computational Linguistics (COLING-88),* Budapest.

Kay, M., J.M. Gawron and P. Norvig (1994) *Verbmobil: A translation system for face-to-face dialog,* CSLI, Stanford.

Landsbergen, J. (1987) 'Isomorphic grammars and their use in the ROSETTA translation system' in M. King (ed.), *Machine Translation Today: The State of the Art,* Edinburgh University Press, pp. 351–372.

Nirenburg, S. and L. Levin (1992) 'Syntax-driven and ontology-driven lexical semantics' in J. Pustejovsky and S. Bergler (ed.), *Lexical semantics and knowledge representation: Proceedings of the first SIGLEX Workshop, Berkeley, 1991,* Springer-Verlag (Lecture notes in Artificial Intelligence), pp. 5–20.

Pollard, C. and I. Sag (1994) *Head-driven phrase structure grammar,* Chicago University Press, Chicago.

Poznanski, V., J.L. Beaven and P. Whitelock (1995) 'An efficient generation algorithm for lexicalist MT', *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics (ACL-95),* Cambridge, Mass..

Sag, I. (1995) 'English relative clause constructions', ms. Stanford University.

Sanfilippo, A., E. J. Briscoe, A. Copestake, M. A. Marti and A. Alonge (1992) 'Translation equivalence and lexicalization in the ACQUILEX LKB', *Proceedings of the 4th International Conference on Theoretical and Methodological Issues in Machine Translation (TMI-92),* Montreal, Canada.

Shieber, S.M. (1993) 'The problem of logical form equivalence', *Computational Linguistics,* **19** (**1**), 179–190.

Talmy, L. (1985) 'Lexicalization patterns: semantic structure in lexical form' in T. Shopen (ed.), *Language typology and syntactic description 3. Grammatical categories and the lexicon,* Cambridge University Press, pp. 57–149.

Trujillo, I.A. (1995) 'Lexicalist Machine Translation of Spatial Prepositions', PhD dissertation, University of Cambridge.

Whitelock, P. (1992) 'Shake-and-bake translation', *Proceedings of the 14th International Conference on Computational Linguistics (COLING-92),* Nantes, France.