SITUATION SEMANTICS AND MACHINE TRANSLATION.

C.J. Rupp
CCL, UMIST
P.O. Box 88
Manchester M60 1QD

Introduction

Situation Semantics is one of the most recent and controversial theories in formal semantics. Machine Translation (MT) is a highly complex application domain, in which research is expensive of both time and resources. On the surface, the space for interaction between these two fields would seem fairly limited, and in practice the application of formal semantics in MT has been very limited, a notable exception being the Rosetta project (Landsbergen 1982, 1987). The abstract translation problem however remains and any application must be based on some formalisation of the problem.

The purpose of this paper is demonstrate that the enriched theoretical vocabulary of Situation Semantics offers a more intuitive characterisation of the translation process, than was possible using more traditional semantic theories. This demonstration will take the form of a formalisation of the most commonly used method for MT in terms of Situation Semantic constructs. In this respect this paper follows on from a previous paper (Johnson, Rosner & Rupp 1988), in which MT was presented as a testing ground for semantic representation languages. This paper will turn the issue around and consider what the theory of Situation Semantics has to offer to an MT application. The abstract description of the MT system to be considered will therefore remain the same.

The paper consists of a basic introduction to the machinery of Situation Semantics, an examination of the problem of translation, a formal description of a transfer-based MT system and some examples of the kind of lexical transfer one would expect to define in such a system.

Situation Semantics: The Basics.

Situation Semantics is an informational approach to formal semantics. The philosophical basis of this theory is laid out in Barwise and Perry's *Situations and Attitudes* (1983) (henceforth B&P). Most of the structure of the theory can be seen as arising out of three basic concepts.

Attunement:

... an organism must be attuned to similarities between situations, what we have called uniformities, and to relationships that obtain between these uniformities... (B&P p10)

Constraints:

... systematic relations of a special sort between different types of situation,... These systematic constraints are what allow one situation to contain information about another. Attunement to these constraints is what allows an agent to pick up information from one situation about another. (B&P p94)

Partiality:

Situation types are partial. They don't say everything there is to say about everyone or even everything about the individuals appearing in the situation type. (B&P p9)

The other main features of the theory can be seen as arising out of the interaction of these three concepts. The combination of attunement with constraints, when applied to the problem of linguistic meaning, leads to the relation theory of meaning. This states that language users are attuned to a particular type of constraint which relates the situation in which an utterance is made with the situation it is about. Put more formally: a sentence ϕ

relates an utterance situation, u , and a described situation, s .

$u[\phi]s$

The notion of efficiency arises out of the interaction of this relation theory of meaning with the notion of partiality. Natural language expressions only carry a certain amount of information and so only partially determine the range of appropriate utterance and described situations. They can therefore be said to be efficient in that they can be used for various purposes. The notion of efficiency implies a clear distinction between meaning and interpretation. It is only possible to arrive at a full interpretation by anchoring the utterance situation and, as a consequence, the described situation to actual situations. The sentence itself carries only meaning.

This theory is sufficiently fine-grained to permit further distinctions within the utterance situation, which contains two differing types of information: the *discourse situation* and *speaker connections*. The discourse situation is that part of the utterance situation concerned with the external facts of the discourse, such as the identity of the speaker and hearer, the temporal and spatial location of the conversation and perhaps even information about the mental states of the speaker and hearer. The discourse situation must be anchored before an interpretation can be determined. Speaker connections are concerned with the linguistic attunement that must be shared by the speaker and hearer for effective communication. The meaning relation can therefore be restated in terms of a discourse situation, d , speaker connections, c , and described situation, s .

$d, c[\phi]s$

The notion of speaker connections assumed in this paper differs slightly from that used in B&P, which was concerned primarily with determining the reference of certain clearly referential phrases, such as proper names and definite descriptions. In this paper it is assumed that most words are in some sense referential although their referents may be complex partial objects; this would seem a natural extension of the

original notion. Speaker connections are therefore the set of culturally specific constraints to which the users of a particular language are attuned in order to permit them to assign meaning to occurrences of its expressions.

Before considering some more recent developments associated with Situation Semantics, it will be useful to sketch some of the distinctions between this theory and more traditional semantic theories, such as Montague Semantics, with particular reference to the implications that these may have for an MT methodology, as in Landsbergen's Rosetta (Landsbergen 1982, 1987). Partiality is the most obvious characteristic of Situation Semantics, when compared to traditional possible world semantics of the Montagovian variety. In traditional theories truth conditions take priority over content. The interpretation of a sentence is the set of possible worlds in which it would be true. Each such world is total and therefore fully determines the answer to any possible question that could be asked about it. Some sentences will be necessarily true and be assigned the set of all possible worlds as an interpretation making them indistinguishable from one another. Others, including all sentences with a necessary truth as a constituent part, form sets of logically equivalent sentences each receiving the same interpretation. This results in a situation where attempts to generate a sentence from its interpretation might result in a sentence with completely different content or the required sentence conjoined with a potentially infinite set of necessary truths. Hence B&P's argument in favour of partial interpretations which contain only a fixed amount of information. This is also one of the reasons why MT systems based on Montague Semantics have been predominantly concerned with derivation rather than representing the interpretation of sentences.

The relation theory of meaning also represents a much greater balance between context and content than more traditional theories, where context is usually limited to the determination of a few indexical terms. Although it has not

yet been adequately explored, the contextual side of the meaning relation does implicitly contain the possibility of representing aspects of the informational structure of texts, which is of essential importance in producing representations for languages such as Japanese or German, where informational structure directly affects syntax. It is not possible to treat such languages in any depth even with derivational techniques, when only truth conditional information can be recorded.

Finally traditional semantic theories assume a static and total interpretation function, which assigns denotations to lexical items. This poses two distinct problems when considering translation. Firstly in the case of words with more than one sense it is not obvious how to decide which denotation to choose. What is required is a more dynamic mechanism which permits the preferred reading to vary according to the context. Secondly there is the implicit assumption that the range of possible denotations is common to both languages concerned, and if we reject this assumption we are faced with the metaphysical problem of constructing appropriate denotations for each language out of an unknown set of primitives, with no philosophical explanation for why this problem arises.

Schematic Representations.

One problem with the original version of Situation Semantics is that it does not have that much to say about the mapping from natural language to Situation Semantic interpretation. The fragments given by B&P are essentially hand coded and give no indication as to how Situation Semantics might be incorporated into a larger more syntactically oriented grammar. More recent work by Fenstad, Halvorsen, Langholm and van Benthem (1987) demonstrates a method of incorporating Situation Semantics into LFG grammars, and HPSG (Pollard & Sag 1987) adopts a similar approach. The combination of unification-based grammar formalisms with Situation Semantics is a very natural move given the role played by partiality in both theories. (See for example Barwise's

comment (Barwise & Perry 1985, p143). Unification-based approaches to Situation Semantics generally require the inclusion of a level of abstract representation which only partially determines the range of possible interpretations. This can be seen as the meaning carried by the sentence under a range of interpretations, but it will be less ambiguous that the original sentence as different syntactic analyses will give rise to different representations.

It is possible to state the meaning relation imposed by such a representation or situation schema in the same terms as were used for the sentence.

$d, c[sit. \phi]s$

The relation between a sentence, ϕ and its representation $sit. \phi$ will be given by a grammar G , which maps strings of a language L to members of a class of representations R , (which will be in the form of Directed Acyclic Graphs). In order to reflect the semantic relation between these two objects it will be necessary to define two auxiliary "interpretation functions" which determine the set of possible interpretations so that

$InL(\phi) = \{ \langle d, c, s \rangle \mid d, c[\phi]s \}$

$InR(sit. \phi) = \{ \langle d, c, s \rangle \mid d, c[sit. \phi]s \}$

The grammar then defines the relation

$G = \{ \langle \phi, sit. \phi \rangle \mid \phi \in L, \\ sit. \phi \in R, \\ InR(sit. \phi) \subseteq InL(\phi) \}$

It could be argued that the introduction of an extra level of representation could pose some problems for the foundations of the theory in that it inevitably attracts comparisons with Discourse Representation Theory (DRT) and representational theories of semantics which assign psychological significance to their intermediate levels of representation. The key to understanding the nature of situation schemata is to see them as containing just the information which may be carried by the use of the construct they represent. Their significance lies

therefore not in the minds of the language users but in the communicative interaction between them. This makes this level of representation the perfect medium for the study of translational equivalences.

Translation Equivalence within a Situational Framework.

This section is concerned with the two essential problems of any approach to MT: the nature and extent of the information that must be preserved, and the nature of the alteration which must be effected. Following on from the previous section it would seem that a partial representation which carried the content of the text ought to supply sufficient information to be preserved. This would represent the meaning of the text while leaving ambiguities of interpretation underspecified. This would effectively freeze the described situation, leaving the context side of the meaning relation as the only domain for translation operations.

A text places fewer constraints on its context than a conversation, because the author and reader know a lot less about each other than do the corresponding speaker and hearer. It follows from this that much of what a text does have to say about its context will remain constant under translation. If an author assumes his reader to have specialised knowledge of a particular subject domain then this requirement should not be affected by translation. This type of information is external to the text and therefore would not appear in the representation of the content and so would not be affected by translation. The only major alteration required in the context of the text is that the reader and author are considered to be users of the target language rather than the source language. This will not greatly affect the external facts of the interaction so the discourse situation can remain constant. It will however drastically affect the linguistic attunements that the author and reader must share in order to communicate. These are culturally conditioned and affect not only the way that words may be used to refer to the uniformities that make up the content of

the text, but also the range of uniformities that it is possible to refer to. This association of linguistic forms with uniformities in the real world is provided by speaker connections and these will be the domain over which translation must operate. Speaker connections do however cover certain text internal forms of reference, such as anaphoric binding; these should also remain predominantly impervious to translation. It is mainly those connections involved in reference into the described situation that must be altered. While this domain only represents a very small part of the situational formalisation of the meaning relation it still represents a vast area of potential variation.

Transfer-based MT.

The problem space for MT is traditionally viewed as being triangular in shape (Vauquois 1979). In this model the problem of translating between texts is reduced to that of a *transfer* mapping between abstract representation of those texts. It is usually assumed that there is a direct relationship between the complexity of the transfer operations and the level of abstraction of the representations; some of the issues involved in this trade-off are discussed in Krauwer & des Tombe (1984). The limit case is where the representations are sufficiently abstract that transfer becomes vacuous: this is exemplified by the interlingual approach adopted in Rosetta (Landsbergen 1982). Increased abstraction can however lead to the loss of relevant information and implies recourse to a culturally independent set of primitives. The adoption of a situational framework for an MT system places interesting constraints on the method to be employed, because both the abstract representational level and the nature of the transfer mapping are determined by the theoretical framework. Interestingly, this turns out to be the kind of transfer-based method most commonly advocated within syntactically oriented approaches to MT.

Within the current model, with situation schemata functioning as the representational level in a transfer-based MT system, the abstraction from text to representation would be that

defined by the grammar relation, G , above, except that two versions of this relation are now required. The parsing relation would be given by a source language grammar, G_{SOURCE} .

$$G_{SOURCE} = \{ \langle \phi, sit.\phi \rangle \mid \phi \in Ls, \\ sit.\phi \in Rs, \\ InRs(sit.\phi) \subseteq InLs(\phi) \}$$

Generation would, similarly, use a target language grammar, G_{TARGET}

$$G_{TARGET} = \{ \langle \phi', sit.\phi' \rangle \mid \phi' \in Lt, \\ sit.\phi' \in Rt, \\ InRt(sit.\phi') \subseteq InLt(\phi') \}$$

The transfer relation can then be defined as a translation relation across representations, TR , expressed in terms of: the two representations, $sit.\phi$ and $sit.\phi'$, a constant described situation, s , and for the purposes of this model a constant discourse situation, d . The actual mapping, K , will be defined across the two sets of speaker connections, c and c' .

$$TR = \{ \langle sit.\phi, sit.\phi' \rangle \mid d, c[sit.\phi]s \\ d, c'[sit.\phi']s \\ K(c, c') \}$$

The translation relation across languages, TL , can then be expressed in terms of the definitions given above.

$$TL = \{ \langle \phi, \phi' \rangle \mid \langle \phi, sit.\phi \rangle \in G_{SOURCE} \\ \langle \phi', sit.\phi' \rangle \in G_{TARGET} \\ \langle sit.\phi, sit.\phi' \rangle \in TR \}$$

In the same way that MT by transfer reduces the translation problem to a translation across representations, so this particular formalisation of the method condenses all the translation operations onto a single K -mapping across speaker connections. This process of restricting the domain over which translation relations hold also reduces their scope. The discussion of translation equivalence was framed in terms of texts, the formalisation of the translation method is expressed in terms of sentence, but speaker connections are a set of constraints on the use of individual words. It might appear that the restrictions of the

transfer mapping to a lexical level smacks of regression towards primitive word-for-word translation, but with the assistance of recent developments within unification-based grammar formalisms nothing could be further from the truth. There are two features shared by UCG (Zeevat et al. 1987), HPSG (Pollard & Sag 1987) and recent versions of LFG (Fenstad et al. 1987, Halvorsen 1987) which make the implementation of such lexical transfer possible. The first is the combination of syntactic, semantic and even phonological information expressed in the same form at all levels of the grammar. This allows for the incremental evaluation of constraints across these various domains. The second is the concentration of information in the lexicon, including information concerning the combinatory behaviour of individual lexical items. These two principles, known as constraint propagation and lexicalism, should make it possible to define lexical transfer relations in terms of the representations associated with individual words of the language, without compromising the ability to specify a wider context.

Lexical Transfer based on Speaker Connections.

Having outlined an approach to translation based on transfer relations over the representations associated with individual lexical items, it is necessary to consider how such an approach might be implemented. This involves two basic issues: the formal nature of such relations and the information that they must express. This discussion will be based on an MT prototype under development at ISSCO Geneva (Johnson et al. 1988) which employs a grammar development tool for unification grammars known as UD, or Unification Device (Johnson & Rosner 1989). Within this environment a representational format has been developed based on the situation schemata of (Fenstad et al. 1987). This will be the framework in which the issue of lexical transfer over graph-structured representations will be considered.

One obvious point of reference in

considering relations between attribute-value graphs is the kind of lexical rule found in PATR type environments (Shieber 1984, Karttunen 1986). These are essentially relations between graphs and are used to treat such phenomena as passivisation. A similar mechanism could be used to implement lexical transfer relations. There would however be one major change in the formulation of such rules, namely the fact that the representations to be related belong to different grammars and so are associated with different syntactic structures. This would affect the way that the root of the graph was associated with the lexical item and the way that information about the surrounding context was passed on. In a lexical rule information to be preserved can simply be equated, but here a translational equivalence is required. There are a number of ways in which this correspondence between elements of different domains might be treated. These include the kind of *structural correspondence* used for relating syntax and semantics in recent work on LFG (Halvorsen 1987, Kaplan & Halvorsen 1988, Kaplan, Netter, Wedekind & Zaenen 1989) and also bilingual lexical entries as in Beaven & Whitelock (1988). The UD formalisation given here will assume a slightly more flexible version of the latter approach, in that not only is the requirement to associate entries of different lexicons recognised, but also the need to be free of the immediate syntactic structure.

Before commenting on a UD implementation of such lexical relations it is necessary to point out that the UD environment does not support lexical rules of the form mentioned above. There is instead a more generalised notion of relational abstraction over the representational domain. This permits the relational characterisation of most of the phenomena usually treated by lexical rules, but not the interpretation under which such rules carry out non-monotonic operations on existing structures. Relational abstractions also permit lexical relations to be broken down into more specific generalisations allowing for a more modular treatment of such phenomena.

Some examples may demonstrate how this technique might be applied to some of the less trivial equivalences between representations. The often quoted equivalence between the Dutch sentence

Ik zwem graag.
I swim willingly.

with the representation
and the English

I like to swim.

involves a difference in the syntactic category used to express essentially the same uniformity. This would be reflected in the structure of the semantic representations assigned to these sentences.

Fig.1 - A Situation Schema for the Dutch sentence: *Ik zwem graag.*

The Dutch representation (Fig.1) shows *graag* as a relation over pairs of relational objects where the English (Fig.2) represents *like* as a relation between an entity and a situation.

Fig.2 - A Situation Schema for the English sentence: *I like to swim.*

The relation between the semantic representations of the two words can be expressed by the abstraction LTR (Lexical Transfer) as follows:

```

LTR(Dutch,English)
  <Dutch ind type> =
  <English ind type> = sit
  <Dutch cond> = [DC]
  <English cond> = [EC]
  <EC rel> = like(_)
  !TR(<DC arg 1>,<EC arg 1>)
  <EC arg 2 ind type> = sit
  <EC arg 2 cond> = [ESC]
  <DC rel ind type> = rel
  <DC rel cond> = [DRC]
  <DRC rel> = graag(English)
  <DRC arg 1> = <DC rel ind val>
  !TR(<DRC arg 2>,<ESC rel>)
  <ESC arg 1> = <EC arg 1>

```

In practice this definition would not require quite so much code as it would be more efficient to draw on abstractions generalising across large numbers of translation relations. The only other abstraction referred to here, TR, is the necessary reference to translational rather than equational equivalence where reference is made to the wider context of the two representations. This

definition is framed solely in terms of the semantic representations and the direct connection between the two representations is made by embedding one under the lexical leaf of the other. This method of representing the correspondence between actual lexical entries is highly experimental and has not yet been applied to any of the larger grammars developed within the UD formalism. It does however avoid one of the more basic problems with the UD implementation of situation schemata: the fact that while it is relatively easy to assert the existence of a piece of representation, it is not possible to ensure that this representation be associated with an actual piece of syntactic structure. It is relatively easy to emulate projection mechanisms such as the σ and ϕ mappings of Halvorsen & Kaplan (1988) by the use of attributes within a larger representation, but it is not currently possible to reproduce the corresponding inverse mappings.

The relation defined above could be described using conventional LFG notations and projections, including a τ mapping for translation between semantic representations. This would however require a slight alteration in the representation language to permit only one condition on each object. The resulting definition would consist of the following set of equations, in which * is the c-structure node associated with the word *graag*, σ a projection from c-structure to semantic representation and τ the transfer projection from source representation to target representation. (In Kaplan et al. (1989) translation relations are predominantly defined in terms of a projection τ across f-structures and the semantic projection is referred to as τ' .)

$(\sigma * \text{ind type}) = (\tau\sigma * \text{ind type})$
 $(\sigma * \text{ind type}) = \text{sit}$
 $(\sigma * \text{cond rel ind type}) = \text{rel}$
 $(\sigma * \text{cond rel cond rel}) = \text{graag}$
 $(\tau\sigma * \text{cond rel}) = \text{like}$
 $\tau(\sigma * \text{arg 1}) = (\tau\sigma * \text{arg 1})$
 $\tau(\sigma * \text{cond rel cond arg 2}) =$
 $(\tau\sigma * \text{cond arg 2 cond rel})$
 $(\tau\sigma * \text{cond arg 2 ind type}) = \text{sit}$
 $(\sigma * \text{cond rel ind type}) = \text{rel}$
 $(\sigma * \text{cond rel cond arg 1}) =$
 $(\sigma * \text{cond rel ind val})$
 $(\tau\sigma * \text{cond arg 2 cond arg 1}) =$
 $(\tau\sigma * \text{cond arg 1})$

This also expresses the translational equivalence of the two words purely in terms of their semantic representations, but this formalism does in principle permit the definition of inverse projections so that $\sigma^{-1}\tau\sigma^*$ would be the c-structure node associated with the word *like*. In order to take advantage of this device it is however necessary to sacrifice the increased expressive power of a representation language defined in UD and the highly modular treatment that the use of relational abstractions provides. While both of these formalisms permit the specification of the required lexical transfer relations without the need to route all information through the source language syntax, it would seem more appropriate to explore the range of appropriate transfer relations within the UD formalism.

A more complex example of this kind of lexical transfer relation might be taken from comparing the use of verbs of motion in English and French. In French the verb often describes uniformities associated with motion and its direction and any specification of manner might have to rely on an adverbial modifier as in

Il descend la rue en courant.
He go-down the road runningly

The representation associated with this sentence is given in Fig.3.

Fig.3 - A Situation Schema for the
 French sentence:
Il descend la rue en courant.

In English however verbs are usually more concerned with the manner of the motion and the direction is left to a prepositional phrase, as in the corresponding

He runs down the road.

which receives the representation in Fig.4.

Fig.4 - A Situation Schema for the English sentence:
He ran down the road.

The following abstraction defines a relation between the English verb *run* and the French verb *courir*.

```
LTR(English,French)
  <English ind type> =
    <French ind type> = sit
  <English cond> = [EC]
  <French cond> = [FC]
  <EC rel> = run(French)
  <FC rel> = courir(_)
  !Trans(<EC arg 1>,<FC arg 1>)
```

```
LTR(English,French)
  <English ind type> =
    <French ind type> = sit
  <English cond> = [EC]
  <French cond> = [FC]
  <EC rel> = run(French)
  <FC rel cond> = [FCR]
  <FCR rel> = courir(_)
  !Trans(<EC arg 1>,<FC arg 1>)
  !Trans(<EC arg 2>,T)
  <EC arg 2 ind type> = loc
```

There two are clauses, denoting two possible transfer relations. The first treats the simple case where there is no directional prepositional phrase in the English and the correspondence is very simple as both verbs refer to a relation over one entity argument. The second clause however treats the case where the second argument to the English verb is locational, as in the example above, which causes the corresponding French expression to be construed as a relation over relational objects, and therefore an adverbial modifier. The formation of such a modifier must be left to the French syntax and the main verb must be supplied by the translation of the prepositional phrase. This definition makes a lot of assumptions about the kind of transfer relations that will be defined on other words. It also implies that a word like the French *en* that makes no contribution to the semantic representation, as it only converts a verb form into an adverbial phrase, should be capable of performing major changes to the translational behaviour of the verb it applies to.

In the previous example the distribution of information across the uniformities that the two languages were able to refer to was clearly different and this was treated by a complex structural relation over the representations assigned to the words involved. With a relatively simple case like this it might seem more appropriate to appeal to some deeper level of semantic primitives to resolve such differences, but the task of lexical decomposition over multilingual vocabularies is obviously doomed by the fact that for most domains there is no obvious constraint on where a culture will decide to draw a distinction. This holds for most domains covered by open

class words that provide most of the content of a text. There are however a few domains that are so structured that they are amenable to decomposition. When languages refer to uniformities in these domains it is usually with constructs that are systematically incorporated into their morphology or with words from distinctly closed classes. These domains correspond to areas that have often caused major problems for MT, such as tense, aspect, modality and determination. In some of these domains it has already become accepted to appeal to an abstract representation that is essentially language independent, as in the work of van Eynde on tense and aspect (e.g. 1988). It is interesting that the primitives required for such representations correspond to the kind of structural relations required in the different object domains of Situation Semantics: locations, relations, situations and entities. It is not possible to present any interesting examples of the treatment of such phenomena here as there is still much work to be done on determining appropriate sets of primitive structural relations, though Cooper (1985,1986) presents a basis for the treatment of tense and aspect within Situation Semantics.

Conclusion.

This paper has presented a formal description of an approach to MT that is based on principles drawn from Situation Semantics, but which utilises the same basic architecture as more syntactically motivated systems. It has also presented some examples of how such an approach might be implemented within current unification grammar formalisms. While this approach has yet to be implemented on any major scale, related work at ISSCO, Geneva has produced grammars for moderately large fragments of German and French which deliver the kind of representation required by such a system.

Acknowledgements.

This research was supported by an SERC studentship (No. 8632138X).

I would also like to express my gratitude to Rod Johnson for providing both intellectual and technical support for this research.

References.

- Barwise, J. and J. Perry 1983. *Situations and Attitudes*. Cambridge, Mass.: MIT Press.
- Barwise, J and J. Perry 1985. Shifting Situations and Shaken Attitudes. *Linguistics and Philosophy* Vol 8 No 1, 105-161. Dordrecht: Reidel.
- Beaven, J.L. and P. Whitelock 1988. Machine Translation Using Isomorphic UCGs. *Proceedings of COLING 88*. Vol 1 32-35, Budapest.
- Cooper, R. 1985. Aspectual Classes in Situation Semantics. Report No. CSLI-85-14 Stanford: CSLI.
- Cooper, R. 1986. Tense and Discourse Location in Situation Semantics. *Linguistics and Philosophy*. Vol. 9 No 1, 17-36. Dordrecht: Reidel.
- Fenstad, J.E., P-K. Halvorsen, T. Langholm and J. van Benthem 1987. *Situations Language and Logic*. Dordrecht: Reidel.
- Halvorsen, P-K. 1987. Situation Semantics and Semantic Interpretation in Constraint-based Grammars. Report No. CSLI-87-101 Stanford: CSLI
- Halvorsen, P-K and R.M. Kaplan 1988. Projection and Semantic Description in Lexical-Functional Grammar. *Proceedings of the International Conference on Fifth Generation Computer Systems, FGCS-88*. Tokyo.
- Johnson, R, M. Rosner and C.J. Rupp 1988. Situation Schemata and Linguistic Representation. Presented at the Lugano Workshop on Computational Linguistics and Formal Semantics, September 1988.

Johnson, R, and M. Rosner 1989. A Rich Environment for Experimentation with Unification Grammars. *Proceedings of the European ACL 1989*. Manchester

Kaplan, R.M., K. Netter, J. Wedekind and A. Zaenen 1989. Translation by Structural Correspondence. *Proceedings of the European ACL 1989*. Manchester

Karttunen, L. 1986. D-PATR: A Development Environment for Unification-Based Grammars. *Proceedings of Coling 86*. 74-80. Bonn.

Krauwier, S. and L. des Tombe 1984. Transfer in a Multilingual MT System. *Proceedings of Coling 84* 464-467. Stanford.

Landsbergen, J. 1982. Machine Translation Based on Logically Montague Grammars. *Proceedings of Coling 82*.

Landsbergen, J. 1987. Isomorphic Grammars and their Use in the Rosetta Translation System. In King, M. (Ed) *Machine Translation Today: the State of the Art. Proceedings of the Third Lugano Tutorial, Lugano, Switzerland, 2-7 April 1984*. Edinburgh University Press.

Pollard, C. and I.A. Sag 1987. *Information-based Syntax and Semantics: Volume 1 Fundamentals*. CSLI Lecture Notes Series No. 13. Stanford: CSLI.

Shieber, S.M. 1984. The Design of a Computer Language for Linguistic Information. *Proceedings of Coling 84* pp 362-366. Stanford.

van Eynde, F. 1988. The Analysis of Tense and Aspect in Eurotra. *Proceedings of Coling 88* Vol 2 pp 699-704. Budapest.

Vauquois, B. 1979. Aspects of Mechanical Translation in 1979. Conference for Japan IBM Scientific Program (GETA Report).

Zeevat, H., E. Klein and J. Calder 1987. Unification Categorical Grammar. In Haddock N, E. Klein and G. Morrill (Eds) *Working Papers in Cognitive Science, Vol. 1: Categorical Grammar, Unification Grammar and Parsing*. University of Edinburgh: Centre for Cognitive Science.