[Proceedings of the *Conference on Theoretical and Methodological Issues in Machine Translation of Natural Languages*, Colgate University, Hamilton, New York, August 14-16, 1985]

Recovering the Speaker's Decisions during Mechanical Translation

David D. McDonald
University of Massachusetts at Amherst

July 1985

1. Abstract

When studied as a source of insight into the human language faculty, rather than to construct a commercially useful service, mechanical translation (MT) is carried out by coupling an otherwise normal natural language parsing system to a normal natural language generation system. In this paper we propose that a crucial capability has been omitted from the design of the parsers that have been used to date, namely a facility for recognizing the information that is implicit in the form of any well written text, matters of emphasis, whether a fact 13 new or old, whether a relationship is given explicitly or left as an obvious inference, signals of intended moves in the discourse, and other things of this sort. We claim that mechanical translations are "mechanical" principally because they pay no attention to information of this sort, and propose that this can be dealt with by incorporating into the parser knowledge of the relationship between usage and form of the sort that is commonplace in any modern language generation system.

2. Why Mechanical Translation as a Task for Study?

The bulk of the research on mechanical translation is motivated by very practical goals. Modern international commerce, especially in the EEC, depends on armies of translators to render contracts, trade agreements, crucial memoranda, etc. into the native languages of the parties concerned. This practical need prompts the designers of most MT systems to continually take what other researchers in natural language processing would consider to be drastic shortcuts in achieving their ends.

The greatest shortcut in commercial MT is the elimination of any conceptual understanding of what the text means. For people, extraction of meaning is an inescapable part of language processing; we truly cannot see or hear a text in our language without knowing at once what it's saying, a process that is so automatic for us that it is arguably part of the encapsulated language module in our minds (Marslen-Wilson & Tylor 1985). Human translators are not behaving any differently when they translate a text from one language to another (though in simultaneous translations they are often too preoccupied to appreciate nuances and finer implications). They are aware of the meanings of the words and context and continually draw on this semantic knowledge in making their judgements about how to phrase the texts in the target language. Imprecision or inaccuracies in target word choice are then usually due to the translator's lack of knowledge of technical terms and idiomatic conventions on the part of, rather than to not understanding what the texts mean.

A commercial MT system cannot be aware of meanings the way a person is for the simple reason that the encyclopedic knowledge base that would be required to capture the requisite range of meanings that occur in commercial texts is well the beyond the state of the art in AI today (Lenat et al. 1984). A fully encyclopedic knowledge base is not required, however, if the MT system is being built only for scientific study: small, topic-specific knowledge bases ("microworlds") that can provide the basis of a system's understanding are constructed regularly today, and the broader AI research community has developed a certain skill in judging whether a "microworld" is in fact well suited to the problems a researcher is using it to study.

What then is the scientific concern in MT—what, if not commercial utility, is the reason to study MT. Languages differ in the communicative devices they provide, especially in their lexical repertoires as is well known, but also in their syntactic devices. As a task domain, MT is well suited as an arena for the study of the match between these linguistic devices that various languages makes available and the different kinds of information people want to communicate. This paper will describe now this can be studied now in MT, while avoiding the ancillary problems that would be part of studying them, AI research on understanding or generation.

3. Translation as parsing followed by generation

From the 1970's onward, researchers who were not attempting to build practical systems have all approached translation in the same way: a conventional language parsing system is used to read the text and construct a suitable internal representation (n.b. the same representation that would have been constructed had the purpose been question answering or cooperating in a task-based dialog). An equally conventional language generation system then takes that internal representation and constructs from it an output text in the target language. This approach has been most visibly used by the AI Group at Yale, the most recently published work is by Lytinen (1984), Isizaki (1983), and Jacobs (1983).

3.1 Information not seen by today's parsers

The approach is sound in its general form. Instantiations of it thus far, however, are flawed by a failure of their parsers to adequately reflect the needs of the generator. Present parsers fail to notice the style of the input text, what it emphasizes, what it marks as a new or unusual event or attribute, what implicatures it invites and how explicitly, or any other information implied by the form of the text beyond its propositional content. All of this information is necessary input to the generator if the target language text is to be a proper reflection of what the speaker of the source language text intended.

It is not surprising that this sort of information is not recovered by today's parsers. The task domains where AI people have most commonly studied language understanding, unlike MT, do not yet place these requirements on parsers for the simple reason that the programs behind the natural language interface, the diagnosis programs, DBMS, etc. would not know how to benefit from them; the information is too subtle to make a difference in their behavior.

In Clippinger & McDonald (1983) we argued that attention to this kind of detail ought to be able to increase the speed and accuracy with which intended inferences could be drawn from a text. We held that well-written texts were typified by a very consistent style with careful attention to consistency in the rhetorical uses that different linguistic constructs are put to. This deliberation on the author's part licenses the audience to read in between the lines of what is literally said to extract information about intended emphasis from where an item appears within the text, or to notice that part of an event is new or unusual by appreciating what others parts might have been mentioned but were not. Attention to this sort of detail should allow a parser to make just those inferences that the author intended without having to do any first-principles reasoning from the literally given information.

In that paper we looked closely at texts like this one, the lead sentence of an April 1982 article in the New York Times.

> "Two Palestinian teenagers were killed and another Arab wounded as the Israeli Army continued to use gunfire to suppress rioting on the West Bank"

This is a very heavily loaded sentence, yet the bulk of the information it conveys comes through only through inferences. Any translation of it that was going to keep the tone of the original would have to do the same.

The major piece of information left to inference is that the Israeli Army did the killing; the sentence only actually says that it occurred "as" the army was putting down riots. If a translation gave the information directly it would completely change the text's force (e.g. imagine what the first line of the equivalent article would have been in an Arab newspaper). The problem for the MT system is to insure that even though the frame that represents the killing event in the internal representation the parser creates has "Israeli Army" in its agent slot, that further annotation is also kept so that the generation system will know not to come out with: "The Israeli Army killed two Palestinian teenagers..."

Other information that is communicated through inference includes the fact that the the use of gunfire in riot suppression is not new and that rioting is not unusual. The Times reporter is writing for an audience that is well informed, they will already know that this particular episode of rioting has been going on for two weeks already, and will be somewhat familiar with how Israel is dealing with it. What is new on this particular day is that for the first time someone was killed. The killing accordingly receives the highest prominence that it can in an English text.

The kinds of reasoning I am arguing that a parser should perform are quite subtle, yet they are what people do regularly, it only reflects the distance we have still to go in our research that the reasoning is at the fringe of the state of the art in AI. I am not aware of any implemented microworlds with well motivated interactive tasks that would benefit from this kind of reasoning, meaning that conventional natural language processing research, based on the development of language-based interfaces to interactive computer programs, will have to wait until one can be developed, which may not be soon.

Mechanical translation research, on the other hand, does not depend upon there being a program that can understand (or motivate) the texts that it is manipulating to prove of effectiveness; instead a person writes the source text and another person judges whether the translation was effective. Accordingly, MT research would let us get on with developing computational models of how texts carry implicit information in their form right now, so that we can begin to modify our parsers and generators to deal with it.

4. Parsing to recover the speaker's decisions

In order to recover this implicit information that the speaker has encoded into the text, the parser needs to continually ask itself the question: why did the speaker of this text give it the form that it has? To make the question tractable, it must be set in the context of the other choices the speaker had available to him: what are they? what parts of the text are they responsible for? what did the speaker achieve by picking the one that leads to this form? Having an inventory of the possible choices will also allow the parser to recognize when an aspect of the text's form is forced by the source language's grammar or by the stylistic conventions the speaker is following; forced choices, of course, do not convey information and can safely be ignored.

The easiest way to answer such questions is for the system to quite literally look at its own generation options and criteria, in effect asking: "if I had said this, what would have been my reasons for giving it the form it has?", and then asserting that those were the speaker's reasons. This will always be a good heuristic evaluation, and will supply an organizing framework for the decisions of the speaker that the system knows it would do differently.

This design restricts the MT system to only noticing form-based inferences that it knows how to generate itself. Moreover it required the system to include a generation capacity for the source language as well as the target language. This does not seem to me to be an unreasonable burden, and it provides an important benefit because one of the things we will want to use MT to understand is just what are the differences in what two languages make available to a speaker, and having both generators side by side will make this much easier to judge.

4.1 Inverting Realization Classes

Specifically I propose that we take the usage-to-form mapping that is already in place in the generator and make it available to the parser.

In our generator MUMBLE, this mapping is given by a data structure we call a "realization class" (McDonald & Pustejovsky 1985b). Formally these classes are functions from concepts or concept types (i.e. terms in MUMBLE's input expressions) to surface structure phrases; the phrases are given as schematic specifications relating fields in the concept (e.g. the arguments to a conceptual relation) to constituent positions within a tree described in terms of grammatical nodes and features, as shown in the second figure below.

A realization class lists the entire set of acceptable realizing phrases for a concept explicitly,[1] and annotates each one to indicate what rhetorical or discourse functions it can serve. These annotations (which we refer to as the "characteristics" of the phrases) are the primary data used by the decision function in selecting a phrase

The figure below shows a realization class we used with the GENARO scene description system (Conklin 1933) as it would be typed into MUMBLE.

```
(define-realization-class LOCATIVE-RELATION (relation arg1 arg2)
   (((Arg1-is-Relation-Arg2)
      ;"The driveway is next to the house."
      clause focus(arg1))
    ((Arg2-has-Arg1-Relation-Arg2)
      ;"The house has a driveway next to it.
      clause focus(arg2) major-minor(arg2, arg1))
    ((There-is-a-Arg1-Relation-arg2)
      ;"There is a driveway next to the house."
      root-clause shifts-focus-to(arg1))
    ((Relation-Arg2-is-Arg1)
      ;"Next to the house is a driveway"
      root-clause shifts-focus-to(arg1) final-position(arg1))
    ((with-Arg1-Relation-Arg2)
      ;"...with a driveway next to it."
      prepp modifier-to(arg2)))
```

A representative Realization Class

Realization classes are associated with concepts by an "assignment" operation, the set of assignments constitute the interface between MUMBLE and the program it is generating for GENARO used an internal representation based on lists, so in this case there would have been an assignment linking any list that appeared in MUMBLE's input whose first element was indicated as representing a locative relation, for example (NEXT-TO DRIVEWAY1 HOUSE1), and this realization class. The

[1] Actually only the alternative wordings and major constituent orderings are explicit; completely predictable transformational variants such as WH-movement, conjunction reductions, or Equi are left implicit and are induced by context as needed. Note that in all of these variants the difference in form is dictated by the grammar and thus cannot be used to convey information.

term NEXT-TO would be bound to the class's "relation" parameter, DRIVEWAY1 to its "arg1" parameter, and HOUSE1 to "arg2".

The phrases that are allowed to realize this relation are given by functions, such as THERE-IS-A-RELATION-ARG2, which will construct the phrase if it is selected; the text each one corresponds to is indicated by the comment just below the function. The annotating characteristics are symbols like ROOT-CLAUSE, and relations over the class's parameters like SHIFTS-FOCUS-TO(ARG1). They act to summarize those aspects of the phrase's form that are germane to decisions about whether it should be used, and allow the encapsulation of irrelevant details of form within the function and away from the decision maker's attention, e.g. required morphological specializations or later syntactic constraints.

Annotating the THERE phrase with the characteristic ROOT-CLAUSE marks it as usable only in certain syntactic contexts (i.e. main clauses). The characteristic SHIFTS-FOCUS-TO(ARG1) indicates that when used it will signal that the focus of the discourse, the object principally under discussion, has changed from whatever it was to be what is bound to ARG1 (e.g. "the driveway").

"Inverting" a realization class amounts to including the recognition of these predefined, parameterized phrases as one of the tasks the MT parser performs. From the phrase specifications the parser can read out a pre-definable mapping to the realization classes that list them. (Nb. once a phrase's parameters have been semantically interpreted, realization class it would have come from can be unambiguously identified.) Once in the realization class, the parser reads out the characteristics that the generator would have had to select had it been producing the phrase itself.

The crucial step in recognizing the phrases is to see in the text the pattern of constituent ordering and fixed words that the phrase specifies. A search pattern should be easily compilable from the phrase's specification in the generator. The figure below shows the specification for the THERE phrase as it would be typed into MUMBLE.

This specification indicates the form of the text to expect in this case it consists of strictly syntactic relationships, in some of the domains that we are using MUMBLE for there are "phrases" at the paragraph level which include rhetorical terms as well

```
(define-choice THERE-IS-A-ARG1-RELATION-ARG2
   :parameters (relation arg1  arg2)
   :phrase (basic-clause ( )
                subject  there
                predicate (VG-NPcomp ( )
                           verb be
                           NPcomp (verbless-nominalized-clause ( )
                                   subject (indefinite)
                                   predicate-adjective (prepp ( ))

   :map ((arg1 . (predicate NPcomp subject))
         (relation . (predicate NPcomp predicate-adjective prep))
         (arg2 . (predicate NPcomp predicate-adjective prep-obj))))
```

A phrase specification

When the MT parser notices that it has parsed the text pattern corresponding to that phrase, it can work back, through the mapping to the realization class, where it will find the two characteristics. The first, ROOT-CLAUSE, will not be of much use to it since it just indicates a feature the form it has just found. The second on the other hand, SHIFTS-FOCUS-TO(ARG1), indicates that the sentence signals a deliberate move in the speaker's (source text's) discourse. This signal, that the object principally under discussion has shifted to the referent of the object of the sentence, must be matched in the output text by the corresponding signal in the target language. The parser must include this "directive" in its representation of the sentence, and not just the sentence's "literal meaning", e.g. that some driveway exists.

4.2 The control structure of the parser is not germane

in considering how a parser could take on these new capabilities, we do not need to be particularly concerned with it control structure will use, present day conceptual analyzers, wasps, ATNs, word-experts, etc. are alike in not attending to this secondary sort of information that choice of form communicates, so no one approach is better than another on those grounds, we should focus instead on the new knowledge sources the parser will need to draw on (e.g. the generator's realization classes and phrase specifications), and on whether any new intermediate representations must be added (see below).

In control terms virtually every AI parsing system of which ! am aware can be rationally redescribed in terms of the three processes: Predict, Scan, and Complete, identified by Earley (1968). The substantive differences in control lie in whether, for example, completion is organized around concepts or around major syntactic categories (e.g. a conceptual analyzer versus as syntax-driven system), or in whether predictions arise bottom up from the lexical entries of words or topdown via recursive descent through the productions of a context-free grammar, (e.g. a word-expert parser versus Earley's algorithm). Of course in the end, the availability of the new knowledge sources or representations may turn out to influence what techniques we choose for the three control processes, but in the beginnings of this research the issue is not important.

5. After the Parser Finishes

I have spoken so far as though the translation process would be straightforward once we have expanded the MT system's parser to also recover the intent-indicating information implicit in the structural patterns of the source text: the parser would render the basic content of the text (i.e. the descriptions of who did what to whom) into a set of objects in the conceptual representation (instantiated object frames, scripts, etc.), and would note the characteristics implied by the patterns as it went; at which point we would simply feed this assemblage of objects and characteristics into a generator with the grammar for the target language and sit back and let the translation appear.

Of course this is not the case. With just this augmented conceptual and rhetorical output from the parser we would be in a position to have a generator reproduce the original text in the same language (in itself an interesting test of completeness), but we would not be prepared for a proper translation, at least not one that would sound fluent to the ear of a native speaker. Rather, a proper translation depends (1) having parsed to an sufficiently abstract level, and (2) on there being an adequate match or equivalence between the linguistic devices of the two languages when it comes to expressing these subtle matters of newness, emphasis, etc.

5.1 Level of Analysis

The abstraction requirement is largely to insure that language-specific, idiomatic conceptualizations are treated properly. As an illustration consider the level one must reach in translating the Russian phrase for "I'm cold" into English. Unlike English, Russian is a case-based language where varied and productive use is made of cases such as dative or instrumental that would not be predicted by an English speaker extrapolating from the surface equivalents of those cases in English. Thus a Russian who, say, had gone out into the snow wearing only a light shirt could express what he feels with the phrase "mnje holodna": the first person dative pronoun followed by the adverbial form of the word for "cold", something that might be literally translated as "to me coldly is". (N.b. Russian copular verbs do not appear in present tense.) In parsing the Russian, one must back off a considerable distance from the linguistic specifics: perhaps representing the type of the text no more concretely than as a description of the speaker's personal state at the time of speech.

All vestiges of language-specific linguistic (particularly syntactic) terms must be left-behind in the output of the parser. If this in not done, the system may over generalize from the structure in the source language ana then sound very "foreign" in the target language, as with "to me coldly is". More seriously we would also imprison our MT system within our local family of languages and be unable to project the internal representation that we recover to target languages with a very syntactic and morphological structure.

So if in fact our input representation to the generator—the amalgam of conceptual terms and intentional characteristics recovered by the parser—has been sufficiently stripped of language specific structure,[2] then the generation of the translated text might well be straightforward, if the generator knows how to describe someone as being

[2] One probably cannot even retain abstract, but still linguistic, predicates !ike predication, or if kept, they must be so stripped of grammatical consequences that they may not be very useful: if we parsed the English "I'm cold" as involving predicating the state "cold" of the speaker, then we would have to strip predication of its otherwise automatic consequence in generation that the thing predicated of goes into the subject, otherwise we would get the Russian "ya holodnje" — roughly '"my body temperature is cold"

cold in the target language (i.e. it has a realization class for such concepts/intentions), then it can certainly do so when told to by the parser. That is, so long as we are willing to live with the limitation on our MT system that it can only say what it already knows how to say when working just as a mono-lingual generator, then we have a good chance at achieving a scientifically useful system.

5.2 Choice of linguistic devices

By a useful system I mean one that would allow us to study the relationships between the linguistic devices of two languages when it comes to expressing these subtle matters of newness, emphasis, etc.—-the second matter on which the quality of the translation will depend. Comparison of the expressive power of two languages comes up only in translation, and thus will only be given a computational account through work on MT.

Some languages may well be more expressive of certain informational nuances than other languages are. English, for example, does not allow us to place a main verb like "killed" in topic position; consequently in our first example the Times reporter did not have the option of indicating unambiguously that it was the fact that for the first time people were killed (rather than wounded) in a West Bank disturbance that was the most significant part of the news. On form grounds alone we have no basis for deciding that the news is not that there were two people killed or that they were teenagers.

The fact that languages vary in which aspects of their form are required as a matter of grammaticality (e.g. that subject must precede verb in English), and which are left free to vary and thus made available to carry information complicates the parser's job. There is no way for it to know whether the author selected a pattern intentionally in order to carry information, or whether he was forced into it by the grammar and then decided to make the most of what he had, e.g., forced to have an initial NP and then using it to carry some content information that could just as well been placed in a less prominent spot.

Much of what contributes to cross-language differences in required form is a difference in grammaticalized semantic perspective, languages vary on which of the attributes of a situation they require to be expressed English insists upon having a subject (i.e. the thing the rest of a clause is predicated of) to the point of using the semantically null "it". Russian allows many subjects to be left implicit (perhaps because information about person is carried in the morphology of the verb), but on the other hand requires every verb to indicate morphologically whether the activity it names is completed or in progress, going further in the case of verbs of motion to distinguish between completed action, a single action in progress, and a habitual action in progress.

This same information can of course be given in English, but not simply by uttering a verb. English speakers are required to be explicit about just what it is that, e.g., makes the action habitual: "I go to work every morning", or about what the situation was in which their action was in progress: "While I was riding my bike yesterday". The requirement to include a certain kind of information will lead a speaker to do so to advantage (as when the Times reporter tells us how many of what sort of people it was who were killed). At the same time a Russian speaker's ability to abbreviate, "pronominalize" if you will, the semantic category of aspect allows him to direct the hearer's attention to other information in the text, without the potentially distracting phrases that would have been required to communicate the summary aspectual facts of the situation when they are also important.

it is unfashionable in this egalitarian age to say that one language is "better" than another at some particular task, yet this seems to be an unavoidable conclusion once we include in the speaker's task the implicit communication of emphasis and newness while staying within accepted stylistic conventions of text length and complexity (a source of information in itself). The experiences of translators of poetry and literature tell us that it is the juggling of the cross-language differences in required and available devices that is the most difficult part of their task. It is incumbent upon people who do research on MT without the practical restrictions imposed by potential commercial application, as experts in computational treatments of mental processes, to attempt to formalize these differences and evaluate their consequences.

6. Promissory notes, Directions, Conclusions

As yet, none of the extensions to the parsing process that I have outlined have been implemented. We have made forays into the problem of "inverting" the linguistic patterns in a realization class, going so far as to develop context free rules from the phrase specifications and using them successfully in an implementation of Earley's algorithm.[3]

However we have not gone the rest of the distance and developed a full language understanding system on that design. There are several reasons for this. One is that we now believe that realization classes as we presently define them do not account for the whole story in the generation of a realistic text: we need to also consider an "attachment process" that combines the output of several classes according to the rhetorical relationships between the conceptual objects involved (see McDonald & Pustejovsky 1985a, 1985b). In parsing terms, this would amount to having a parser for a TAG grammar; what impact this will have on the conventional repertoire of control structures for parsing is not yet clear.

6.1 Structured input to the generator

The second reason is potentially more significant and it reflects a promissory note in the proposal of this paper that could turn out to be hard to collect. It is that we do not yet understand enough about the constraints on the structure of the input to the generator. MUMBLE is driven not by a feature space of characteristics and a list of conceptual objects, but by a highly organized "specification" of the sort shown below. The MT parser will have to build such a structure, working from the objects and characteristics, if it is to gain maximal leverage from a generator like ours.

[3] The goal of that work was a tool for automating the construction of the concept-class assignments that MUMBLE uses as its interface to the programs it is talking for. At the time of this writing this tool is not yet complete. It is waiting on the implementation of revisions that we have designed to the notation for MUMBLE's production of function words, and to the notation that defines the relationship between lexical heads (especially verbs) and the transformational families that accompany their subcategorization frames.

At the time of this writing we have neither an underlying program and knowledge base nor a text planner that is sophisticated enough to motivate the kinds of texts this paper has discussed, which has made drawing conclusions and sometimes even conducting suitable experiments very difficult, however we have proceeded by "reverse-engineering" texts by hand and considering what control over the linguistic levels of the generation process a conceptual-level text planner would want to exert. MT experiments would give us a much more empirical basis for judging such experiments since the criteria for success—mimicking the information content of the source text—would be clearer, and the question of whether we had abstracted sufficiently from linguistic forms in our message level would be plainly answerable.

```
(the-day's-events-in-the-Gulf-Tanker-War
    :events-require-clarification-as-to-source
    (main-event #<same-event-type_varying-patient
                    #<hit-by-missiles Thorshavet>
                    #<hit-by-missiles  Liberian> >
                 :unusual #<number-of-ships-hit 2>
                 :identify-the-ships)
    (particulars
        #<damage-report Thorshavet Oslo-officials>
        #<damage-report Liberian Lloyds> ))
```

Corresponding to:

> "Two oil tankers, the Norwegian-owned Thorshavet and a Liberian-registered vessel, were reported to have been hit by missiles Friday in the Gulf. The Thorshavet was ablaze and under tow to Bahrain, officials in Oslo said. Lloyds reported that two crewmen were injured on the Liberian ship."

This "realization specification" (discussed in McDonald & Pustejovsky 1985a) was written by hand as part of an experiment in what a text planner could provide as the proximal source of a complex, deliberately structured text. !t seems to us to tie down a sufficient number of the "expressive degrees of freedom" associated with the realization of conceptual objects as rich as these (indicated by the "#<...>" notation] to force the construction of the text we were trying to match, rather than one with a different ordering of the phrases or an alternative mixture of express and implicit information.

6.2 Directions

What we do not yet know is whether we are presuming a plausible set of capabilities and constraints on the actions of the text planner that would come up with this specification, and therefore whether we are presuming a reasonable structure for input to the generator and, by implication, a reasonable conceptual level target for the MT parser. Studies to establish whether this is the case are more studies of generation than MT per se, however it is high time that they were carried out in multiple languages at once (as has been done at Yale for some time, though with a markedly simpler generation system; see Lytinen (1984), who was working with a generator developed by Rod McGuire). Multi-lingual studies make plain any hidden language-specific presumptions that are implicit in the design of the generator and text planner.[4]

A. potentially high payoff possibility that could emerge from multi-lingual generation studies would be an independent set of criteria for which aspects of the generator's design could be profitably hypothesized as part of the "Universal Grammar" putatively underlying all human competence in language. If the design of generator is otherwise carefully done (from the perspective of the theoretical linguist), then the division between those capacities and constraints provided by its functional architecture and those that come from explicitly represented rules and data that it uses will be significant. A natural hypothesis to entertain is that the functional architecture alone should be responsible for those aspects of the generator's design that can be shared across languages; knowledge that is only embodied in an explicit representation that the functional architecture manipulates should vary a great deal more, and would thus be a natural candidate for our language specific knowledge.

[4] In the fall of 1960, Edward Hoenkamp developed a small surface grammar of Dutch for MUMBLE and used it to produce Dutch versions of some simple descriptions of Shakespearian plays. Even though only a few weeks were spent on this project, it was enough to determine that MUMBLE's design incorporated a tacit presumption about the placement of prepositions that was inappropriate for SOV languages.

7. Acknowledgements

The ideas presented in this paper were developed over the last several years, and have profited from discussions I have had with Kevin Gallagher and James Pustejovsky, who were also responsible for the context free rule treatment of phrase specifications. Preparation of this paper was supported in part by Darpa contract no. N0014-85-K-0017.

8. References

Clippinger, J. & McDonald, D.D. (1983) "Why Good Writing is Easier to Understand", Proceedings of IJCAI-83, Wm. Kaufmann Inc., pp. 730-732.

Earley, J. (1968) "An Efficient Context-Free Parsing Algorithm" CACM 13, (2) pp. 94-102.

Isizaki, S. (1983) "Generation of Japanese Sentences from Conceptual Representations", Proceedings of IJCAI-83, Wm. Kaufmann Inc., pp.613-615.

Jacobs, P. (1983) "Generation in a Natural Language Interface", Proceedings of IJCAI-83, Wm. Kaufmann Inc., pp. 610-612.

Lenat, B., Borning, A., McDonald, D., Taylor, C, Weyer, S. (1983) "Building Expert Systems with Encyclopedic Knowledge", Proceedings of IJCAI-83, Wm. Kaufmann Inc., pp. 167-169.

Lytinen, S. (1984) "The Organization of Knowledge in a Multi-lingual, Integrated Parser" Yale AI Group technical report #340.

Marslen-Wilson, W. & Tyler, L. (1985) paper presented at UMass Cognitive Science Project's workshop on "Modularity", Hampshire College, June 6-8, 1985.

McDonald, D.D. & Pustejovsky, J. (!985a) "TAG'S as a Grammatical Formalism for Generation", Proceedings of the Annual Meeting of the Association for Computational Linguistics, University of Chicago, July 8-12, 1985.

McDonald, D.D. & Pustejovsky, J. (I985b) "Description-Directed Natural Language Generation" Proceedings of IJCAI-85, Wm. Kaufmann Inc.