

**On the Role of Old Information in Generating Readable Text:
A Psychological and Computational Definition
Of "Old" and "New" Information in the NOSVO System**

Mark Vincent LaPolla
Linguistics Dept.
University of Texas, at Austin
Austin, Tx 78712
CS.LAPOLLA@R20.UTEXAS.EDU

NOSVO is a Natural Language Generation postprocessor which is sensitive to old/new information contrasts. We believe that generating old information first establishes cohesion in text promoting readability. This paper describes the NOSVO system in detail and the motivations for building it. We also provide a psychological and computational definition of "old" and "new" information.

There are situations where the speaker is constrained by a grammatical rule, and there are situations where he chooses according to his meaning ... ; but there are no situations in the system where "it makes no difference" which way you go. This is just another way of saying that every contrast a language permits to survive is relevant, some time or other. (Bolinger 1972:71).

1.0 Introduction

There are at least two stages of text generation. One is generating the content of the text. The other is generating the language that represents and communicates the content (Thompson 1977). These two stages, though interrelated, have their own sets of interesting problems and principles. The first stage, generating the semantic content of the text, involves motivating, planning and creating the conceptual and semantic content of a piece of text. Once the semantic representation for a text has been constructed the language of that text can be generated. The second stage, language generation, involves communicating the intent and content of the text without confusing or misleading the reader. This paper will address the second stage only.

It is not enough to merely generate text. It is also necessary to generate cohesive text. However a shopping list is cohesive, though not "flowing" text by any means. A set of sentences that are propositionally related are cohesive though are not necessarily beautiful prose.

It is not enough to attend to ellipsis and pronominalization to generate readable prose. We believe that there are other factors which must be attended to to generate prose. The NOSVO system is an attempt to take into account old/new information contrasts (Chafe 1974, 1976) which we believe will help natural language generation systems produce more readable text.

2.0 "Old" and "New"

It is important at this point to clarify the use of the terms "old information" and "new information". The term "old information" is a misnomer, though it expresses the intuition needed for this paper. The term suggests "it is what the listener is expected to know already" (Haviland and Clark 1974). The term "new information" is also incorrect. It suggests that what the speaker has uttered is completely new to the hearer and is being introduced into the speaker's consciousness for the very first time. But as Chafe (1976) points out a person uttering the sentence "I saw you father yesterday" is "unlikely to assume that the addressee had no previous knowledge of his father, even though by the usual criteria "your father" would be considered new information" (Chafe 1976 : 30). Chafe suggests that the terms "already activated" and "newly activated" would be more appropriate.

Certainly "already activated" and "newly activated" (information) are by far better than the terms "given" and "new". Even so, they are still somewhat imprecise. "already activated" and "newly activated" imply, as does Chafe, that the concept that is activated, whether extralinguistic or linguistic, is directly activated by a (linguistic or extralinguistic) referent. As Prince (1978) shows for clefts and as LaPolla (1986) shows for inversion and prepositional phrase fronting this need not be the case. Rather the antecedent "simply has to be appropriate to the situation, and hence cooperatively

'assumable' as being there" (Prince 1978:889). This is quite important and expands upon Chafe since for him the important thing is that the antecedent must be in the hearer's consciousness, i.e. in the hearer's focus of attention, while for Prince and LaPolla it need only be appropriate to the situation or in some other way cooperatively assumable to be in the hearer's consciousness.

Hajicová and Vbřová (1981) also takes exception with the terms "given (or old)" or "new" information and suggests that "contextually bound" and "contextually non-bound" lexical item would be more appropriate. "contextually bound" and "contextually non-bound" is even more appropriate than "already activated" and "newly activated" because it seems to also convey situational appropriateness. However, it seems that Hajicová restricts her terminology, as well as her theory of discourse (focus) structure, to linguistic antecedents. That is, her "shared stock of knowledge" appears to be closer to, if not completely, linguistic in representation. Therefore, neither her theory or terminology has the power to deal with an antecedent that is merely inferable or appropriate to a situation. We will use the familiar terms "new/old information" but will define them a little more precisely later in the paper.

3.0 The System

NOSVO (Not Only Subject Verb Object) is a preprocessor that aids in the generation of English text. It is not in and of itself a text generator but does contain a very simple predicate to English translator (similar in spirit, though not in complexity, to Simmons 1984). NOSVO is sensitive to the old/new (i.e. given/new) contrasts in a discourse (Chafe 1974, 1976; Hajicová and Vbřová 1981) and the syntactic structures that allow a writer, or speaker, to begin a sentence with old information. NOSVO organizes the semantic content of predicates to produce an old information first ordering. (Appendix A contains a short example of text with and without the application of the old information first principle.)

The rest of this section describes NOSVO and its motivation in detail. (See Appendix B for a high level data flow diagram of the system.)

3.1 Motivation

One of the goals of communication is to modify and extend the store of information in memory (Hajicová and Vbřová 1981) and language facilitates communication. If the above statement is accepted then one might characterize a discourse as a session where particular slices of memory are accessed and changed. Since this accumulated store of information in memory is presumably very large and has a complicated structure, it is useful if the communicator can first identify the locations in memory that are to be modified or added to before introducing the modifications

or additions (Hajicová and Vbrová 1981; LaPolla 1986).

Beginning a sentence with old information allows a speaker to direct the hearer's focus of attention to exactly those elements which he wishes to modify before he utters the rest of the sentence, presumably new information (Chafe 1974). If a speaker began a sentence with new information, the listener would not know where to connect it in the discourse. The listener would have to wait until the old information is uttered to locate where the new information is to be used (LaPolla 1986). The former process takes less concentration by the listener (Green 1980).

If we model human memory, specifically lexical access, as a spreading-activation network (Quillian 1962, 1967; Collins and Loftus 1975) we can describe the effects of uttering old information before new as priming, for at least the class of *linguistic antecedents*. Old information primes an already active node in memory raising its level of activation and making it more accessible than surrounding nodes. The priming may then spread to related concepts. This causes only the relevant and related portions of memory to be available for modification or replacement, and restricts the amount of memory brought into the processing of discourse. (Our model can be extended to cover nonlinguistic and inferable antecedents by using a hybrid representation (Vilain 1985) [cf. Haviland and Clark (1974) or Prince (1978) for a definition of inferable antecedents].) If new information were uttered first then the level of activation achieved would be equal or close to previously primed concepts and the speaker's attention would not be properly focused.

We can now define old (linguistic) information as already activated memory. New (linguistic) information, then, is either information not in memory or information that has not yet been activated. Using Quillian's model, all concepts that have been primed by the discourse, either directly or by spreading-activation, are old. Everything else is new.

Structures like inversion and PP-fronting, to a name a few, extend the syntactic and logical possibilities of information presentation in a language. They allow adjuncts, e.g. in PP-fronting, or objects, e.g. in passives, or other arguments, e.g. in inversion, to be presented before other syntactic constituents.

To complete the above definitions we will adopt and expand Quillian's definition of "concept". For Quillian, and Collins and Loftus, a concept corresponds "to particular senses of words or phrases" (Collins and Loftus 1975: 408). We will expand this to include extralinguistic phrases or groups of actions and objects. A concept then is not only senses or words or phrases such as NPs, "machine", VPs, "to machine", and phrases "the particular old car I own" (page 408) but also extralinguistic objects: a picture of a fire engine, and actions, the picture of a fire engine racing down the street. This also includes situations such as eating in a restaurant and paying the check. In sum, concepts are linguistic words and phrases as well as more difficult things to pin down like situational scripts.

We have adopted a well proven model of human memory to describe the effects of linguistic information on memory. We have used this theory to model discourse processing as a type of memory processing. We have also advanced the idea that language in its role as a communications facilitator allows a speaker to direct which concepts will be primed and therefore what the hearer will focus his attention on. We have defined old and new information within this framework and have briefly described the contribution inversion, PP-fronting and passivization make to the presentation of information in language. In the rest of this paper we will show how we have integrated these ideas and theories into NOSVO.

3.2 System Overview

NOSVO takes as input a syntactico-semantic predicate representation of each sentence in the text to be produced (following Simmons 1984). (See Figure 1 for an example. The asterisks indicates a backward pointing arc (Simmons 1984).) It determines which of the predicates constituents are "old" information. It then updates its lexical memory to reflect the predicates utterance on the listener's own memory. The lexicon is a semantic network following

Quillian. Currently it supports the relations ISA, HAS, HAS-PART, LOC, IS-CALLED, EXAMPLE, SUPER- and SUBCLASS.

NOSVO assumes that each of the underlying predicates maps to a sentence. This assumption has allowed us to focus on the old/new information distinction that is the core of the NOSVO system without worrying about the mapping from "deep" semantic structure to surface structure. In other words, NOSVO is not a robust English generator and has all of its intelligence devoted to the manipulation of information. We feel that language generation is a difficult research issue and is beyond the scope of this paper and the current system.

```
(spend tns past ae ((father number sing)
agt* (accompany tns pres infl inf)
ae (me number sing))
before* ((be tns past number sing)
atr (tall) atr (enough)
agt (I number sing)
agt* ((ride tns pres infl inf)
on* ((coaster number sing)
atr (big)
ae (me number sing))))
agt ((hour number pl) atr (pleasant) atr (many)))
```

Figure 1

3.2.1 NOSVO

First, NOSVO determines which parts of the input predicate are "old information". It checks the nodes with the five (5) highest levels of activation in its knowledge base (KB), e.g. levels 6-10 where 0 is not activated and 10 is fully activated. If any of the heads of any of the arguments from the input predicate is among the activated nodes then that argument is marked as "old". (NB: We do not address in this paper internal ordering of constituents other than sentences. We only check the heads of constituents directly below the sentence level.)

After NOSVO has marked the old information in the input predicate, it updates the status of its KB to reflect the hypothesized change the generated sentence will have on the listener. To do this, NOSVO first parses each input predicate into its constituents. For each argument in the input predicate NOSVO primes, or reprimed, a corresponding (concept) node in its KBs. (Note: Though we do not check each part of each constituent when determining what is old, every part of every constituent does affect the state of the lexical memory of NOSVO.)

When a node in the knowledge base is primed, it is tagged with a level of activation. The initially primed concept is tagged with the highest level of activation, call it level 10. The initially primed concept is also tagged as the initially primed node, i.e. the node primed by the discourse and not by spreading-activation. Activation then spreads outward raising the level of activation of surrounding nodes. As the spreading-activation gets further away from the initially primed node its effect is reduced proportionally to the distance traveled (following Collins and Loftus 1975). For example, at the initially primed node the level of activation is 10, at the next node it is 9, at the next 8 and so on. NOSVO also tags the surrounding nodes with the name of the initially primed node. We realize that we do not know by what exact proportion the activation effect is diminished. Nor do we know how long activation lasts or at what rate is deteriorates. These are questions for future research.

3.2.2 Generating Sentences

After NOSVO has marked the arguments in the predicate, the result is passed to a simple English generator. The role of the generator is complex but the way it executes its role is simple. The generator looks at the marked predicate and chooses the correct English syntax to map the predicate to English. We realize that this task is very complex and our treatment of it is superficial. We realize that entire systems have been created to address the problem, e.g. MUMBLE (McDonald 1980; see also McDonald, Meteer and Pustejovsky 1987). We also realize that we do not address pragmatic considerations in the

generation of discourse (Hovy 1987). However, recall that our goal was not to create a robust language generator. Our goal was to create a system that could recognize old information in a phrase being generated. After this is done it is the responsibility of the rest of the system to act on that information.

We assumed from the beginning that the underlying semantic representations NOSVO processes are already organized into sentences. This was not to aid NOSVO in its task, though it does by explicitly defining the relationship between the verb and its arguments, but rather to aid the simple predicate to English generator.

The English generator coupled with NOSVO takes NOSVO's output and analyzes it. If a prepositional phrase (PP) adjunct has been marked as "old" then the generator fronts it, e.g. "In the park, John kissed Mary". If an object has been marked as "old" the generator generates a passive, e.g. "The apples were eaten by Vincent". If a pp argument or adjunct of an intransitive is marked as "old" the generator generates an inverted sentence, e.g. "Around the bend lives John". If no old information is found other than in the subject then a simple Subject Object Verb sentence is generated, e.g. "Vinnie loves Mark". If a predicate contains no old information, either explicitly or implicitly, it is a nonsequeter and should not be generated.

At this stage in the generators development only simple syntaxes are generated. Extraposition out of clauses is not addressed. This was not our intent. Also it was not our intent to argue here that the presentation of old information first is the sole discourse function of structures like inversion, pp-fronting and passives but only one of their discourse functions, perhaps the main one. It was our intent to build a system that could determine which parts of a sentence under generation were old information. It was also our intent to clarify the terms old and new information and to put their definitions in perspective both linguistically and psychologically. These issues we have addressed.

4.0 The System In Detail

In this section, we will present the NOSVO system in detail.

4.1 Detailed Overview

NOSVO's grammar is segmented into two parts: i) plain vanilla SVO rules, e.g. S -> NP verb PP, "A little angel stood outside" (Green 1980, page 582), "Uncle Jack lives around the bend" (LaPolla forth coming), and ii) so called old information first syntax, e.g. S -> PP verb NP, "Outside stood a little angel" (Green 1980, page 582), "Around the bend lives Uncle Jack" (LaPolla 1986 and 1988). (Cf Green 1980 and LaPolla 1986 for details). When NOSVO encounters old information in any constituent in a predicate, except for the logical subject, it uses the old information first syntax grammar to generate the sentence. Otherwise it uses the plain vanilla grammar. NOSVO does not do any extraposition from within embedded clauses nor does it handle the differences between internal arguments and adjuncts. It only produces the three variations on standard, plain vanilla syntax discussed above. These issues will be addressed in future versions of NOSVO.

NOSVO has two kinds of knowledge bases from which to work: linguistic and conceptual. There are two linguistic knowledge bases: the lexicon and the discourse base. The lexicon maps into either a domain specific or a non-domain specific KB. Both the domain specific and non-domain specific KBs are hierarchical networks which support the relations ISA, HAS, HAS-PART, LOC, IS-CALLED, EXAMPLE, SUPER- and SUBCLASS. The domain specific and non-domain specific knowledge bases are the Quillian style semantic networks. In NOSVO's current avatar, there is a one to one mapping of lexical items to concepts in the knowledge bases. That is, there is no lexical or conceptual ambiguity. A more robust system would allow for multiple mappings in both direction because of the power and depth it gives. A more robust mapping would produce two problems though. The first would be the extra computation and heuristics required to resolve the ambiguities. The second problem would be determining when to prime a node. It might not be correct

to prime a node just because the lexicon accessed it. One might have to wait for a completed parse before priming the concept bases.

The discourse base is a tree. As NOSVO generates text, it builds the discourse tree connecting old information to new while retaining the autonomy of each predicate. The discourse base contains the structure of the discourse and is a way to record priming. The discourse base maps directly into the KBs as well. One could assume that the discourse is just a section of primed memory. However, it was felt that a more linguistic representation would be useful in helping to resolve anaphora. The discourse base was modeled after the discourse mechanism in LaPolla (1986).

NOSVO's first step in determining whether a predicate or an argument is "old information" is whether or not it has been introduced into the discourse, that is whether or not it is definite (Heim 1982,1983). If a referent has already been introduced into the discourse then it is necessarily old information. However, the converse is not necessarily true. That is, just because a referent has not been introduced into the discourse does not mean it is 'new information' (Von Stechow 1981). All it means is that the referent has not been introduced. The information may have been.

The nonlinguistic KBs contain metalinguistic knowledge about articles, stories or other appropriate formats and the expectations speakers have about them, knowledge about the topic of the text and specific and general knowledge about the lexicon.

NOSVO tries to establish a link from the sentence under generation to one of the KBs. If a link can be found from an argument or an adjunct in the input predicate to the knowledge base, the link is recorded and the old information first grammar is used to generate the argument or adjunct first, if possible. (This is a oversimplification and will be expanded upon.)

If NOSVO can not establish a link to the knowledge base, it searches the meta-knowledge base, i.e. the knowledge base containing information about author motivation for writing an article or story, the techniques authors use to write articles or stories, and information about articles and stories, their parts and subcategories. The meta-knowledge base is primarily used to establish bridges (Clark and Haviland 1974) discourse initially from the old information in the predicate to an inferrable metalinguistic antecedent. For example, a college professor may begin a lecture (or course) with the discourse initial utterance "What we're going to look at today (this term) is..." but not "*What one of my colleagues said this morning was..." (Prince 1978, page 889) or "*What I told my wife this morning was..." (The asterix means semantically unacceptable) The first sentence is allowable because the context, i.e. the class room setting, allows a direct inference to studying (for the term).

The meta-knowledge base has two parts: a taxonomy similar to the Domain Specific KB and scripts that have knowledge about objects and the actions that they perform, e.g. writers write stories, writers set scenes, stories have scenes.

The algorithm which NOSVO uses to determine which old information first syntax is appropriate is straightforward. The complicated part is how NOSVO decides what is old information. Currently NOSVO searches through all concepts with activation greater than 5. The value 5 is arbitrary, however. It is still an open research question when an antecedent can be considered no longer in the speaker's/listener's common ground (Chafe 1974, 1976), or no longer cooperatively assumable (Prince 1978).

4.2 The Components of the System

This subsection will outline in detail the various components of NOSVO and their function. This subsection is organized in parallel with the data flow diagram in Appendix B, starting with NOSVO's first component subsystem.

4.2.1 The Predicate Parser

The Predicate Parser identifies and parses the input

sentential predicate into its component parts. This is the first stage in identifying old information in a predicate.

4.2.2 Predicate Argument Translator

The Predicate Argument Translator translates the linguistic representation of the input predicate constituents into tokens, from the lexicon, which map into the discourse base and the other KBs. Notice that the representation of discourse referents and concepts need not be the same, only that each referent or concept be indexed, and indexable, by a token. The tokens are only used to query the knowledge bases. When we speak of finding a link between the input predicate and the knowledge base that link is established through the conceptual translation.

4.2.3 The Discourse Base Searcher

The Discourse Base Searcher searches the discourse to determine whether any of the input predicate arguments in the predicate have been previously introduced into the discourse. If an antecedent(s) is found the link is recorded and the whole predicate, with highlighted old information, is sent to the Linguistic Converter and Category Analyzer (L.C.-CA).

4.2.4 The Domain Specific KB Searcher

If no antecedents are found in the discourse base the Domain Specific KB Searcher searches the domain specific KB for a possible link.

First maximally primed nodes are investigated, i.e., nodes with priming 10. Then other less primed nodes are investigated and so on up to priming level 5. Note: that the amount of search necessary increases as the priming decreases. If a link is found to a node, that node is primed and the input predicate is sent to the Linguistic Converter and Category Analyzer with the old information highlighted. If not, control and the input predicate is passed to the Non-Domain Specific KB Searcher.

4.2.5 The Non-Domain Specific KB Searcher

The Non-Domain Specific KB Searcher searches for an antecedent in the Non-Domain Specific KB. In NOSVOs case non-domain specific knowledge is general and prototypical knowledge. So, for example, if the domain specific KB is Navy ships, then the *non*-domain specific KB might contain information about ships in general, water, vehicles, transportation, or guns and fighting in general. The exact same mechanism is used to search the Non-Domain Specific KB as the Domain Specific KB. If no antecedent is found in this knowledge base the predicate is passed to the Bridge Building Inference Engine.

4.2.6 The Bridge Building Inference Engine

If all the other processes have failed to find a link from predicate to common ground, i.e. the context, both linguistic and nonlinguistic, of a discourse, NOSVO tries to build a bridge, an inference, which connects information in the predicate to a metalinguistically inferable antecedent. This component of NOSVO is not very robust. NOSVO will eventually be reimplemented using a Valain (1985) type hybrid approach. Then the Bridge Building Inference Engine will be expanded.

At this point, the careful reader may think that given the nature of NOSVO's search mechanisms that it must always succeed in establishing a connection from input to knowledge base. That is the case. Indeed it must be the case. Consider that everything that people say to each other must in some way link to the common ground in order to be understood. Or else the utterance would be a nonsequitur. Even the self introductions performed by two people who do not know each other, and have just met, are expected and reasonable. The formula has a metalinguistic antecedent in culture.

The question for NOSVO is not whether an antecedent exists but rather what it is. If NOSVO cannot find an antecedent it assumes that one exists and generates a sentence with plan vanilla, SVO structure, leaving it to the

reader to establish the connection. If NOSVO did not assume an antecedent it would have to discard the sentence as a nonsequitur and potentially confusing and/or misleading. This is an important issue and will be dealt with in a later and expanded version of NOSVO.

4.2.7 The LC-CA

The Linguistic Converter and Category Analyzer analyzes the old information to determine its syntactico-semantic category. It checks whether it is a prepositional phrase, agent, theme or instrument, in that order. It then decides if the old information is an internal or external argument or a prepositional phrase adjunct. With this information it picks the type of grammar that will place the particular argument or adjunct first and sends the choice along with the predicate to the English generator.

4.2.8 The English generator

The English generator is a prolog grammar segmented into the various old information first syntaxes, e.g. prepositional phrase first, object first rules, and a plain vanilla syntax. At this point all, or most, of the intelligent work has been done and the generator is nothing more than a syntactic manipulator under the direction of the Linguistic Converter and Category Analyzer.

5.0 Future Research and Directions

Future research and development topics for NOSVO include:

1. determining when information cannot be assumed to be in the listener's common ground, i.e., at what level of priming is a concept not in the listener's common ground?;
2. expanding NOSVO's capability to handle ellipsis, definiteness, and pronominalization and investigate how the generation of ellipsis and definiteness affects the generation of old information first;
3. extending NOSVO to do more of the linguistic generation from either a more "conceptual" representation or to take as input another source language such as another natural language or a computer program and generate English from that underlying representation, i.e. expand NOSVO's backend;
4. extending NOSVO's capabilities to handle the subtle distinction between arguments and adjuncts;
5. determining how much the nonapplication or missapplication of the old information first principle, discussed above, makes a difference in reading and understanding text;
6. finally, investigating other old information first syntactic structures and phenomena to determine how they affect a discourse and how they might be integrated into NOSVO.

The next generation of NOSVO will be written in CLOS and Lisp. The application will be "generating descriptions of Lisp programs". CLOS objects will be used to organize the knowledge structures and CLOS methods will be used to do the actual parsing. Eventually NOSVO will be expanded and refined along the directions stated above.

Acknowledgements

I would like to thank Kent Bimson, Mirjam Fried, Randy LaPolla, Marie Meteer and Varda Shaked for all their help and criticism on this abstract. Any mistakes are my own.

References

- Bolinger, D. *That's that*. The Hague: Mouton, 1972.
- Chafe, Wallace. Language and Consciousness. *Language*, 1974, 50(1), 111-133.
- Chafe, Wallace. "Givenness, Contrastiveness, Definiteness, Subjects, Topics and Point of View". In Charles Li (Ed.), *Subject and Topic*. (ed Li, C. N.), New York: Academic Press, 1976.
- Collins, Alan M. and Elizabeth F. Loftus (1975) "A Spreading-Activation Theory of Semantic Processing". *Psychological Review*, 87, 407-428.
- Davison, Alice. "Peculiar Passives". *Language*, March 1980, 56(1), 42-66.
- Green, Georgia M. "Some Wherefores of English Inversions". *Language*, 1980, 56(3), 582-602.
- Hajicová, Eva and Jarka Vrbová. "On the Salience of the Elements of the Stock of Shared Knowledge". *Folia Linguistica*, 1981, 15, 291-303.
- Haviland, Susan E. and Herbert Clark. "What's New? Acquiring New Information as a Process in Comprehension". *Journal of Verbal Learning and Verbal Behavior*, 1974, 13, 512-538.
- Heim, Irene R. *The Semantics of Definite and Indefinite Noun Phrases*. Doctoral dissertation, University of Massachusetts at Amherst, September 1982.
- Heim, Irene R. "File Change Semantics and the Familiarity Theory of Definiteness". In Rainer Bäuerle, Christoph Schwarze and Arnim von Stechow (eds.) *Meaning, Use, and Interpretation of Language*. Walter de Gruyter, Berlin, 1983, 164-189.
- Hovy, E. H. *Generating Natural Language Under Pragmatic Constraints*. Unpublished Yale Dissertation, YALEU/CSD/RR #521, 1987.
- LaPolla, Mark Vincent. "The Role of Inversion, Clefting and PP-Fronting in Relating Discourse Elements: Some Implications for Cognitive and Computational Models of Natural Language Processing". In *Proceedings from the XI International Conference on Computational Linguistics*, 1986, 168-173.
- McDonald, David D. *Natural Language Production as a process of Decision-making under Constraints*, unpublished Ph.D. Dissertation, MIT, Artificial Intelligence Laboratory, 1980.
- McDonald, David D., Marie W. Meteer (Vaughan) and James D. Pustejovsky. "Factors Contributing to Efficiency in Natural Language Generation". In G. Kempen (ed.) *Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology and Linguistics*, Kluwer Academic Publishers, Boston/Dordrecht 1987.
- Prince, Ellen F. "A Comparison of Wh-Clefts and It-Clefts in Discourse". *Language*, 1978, 54(4), 883-905.
- Quillian, M. R. (1962) "A Revised Design for an Understanding Machine", *Mechanical Translation*, 7, 17-29.
- Quillian, M. R. (1967) "The Teachable Language Comprehender: A simulation program and theory of language", *Communications of the ACM*, 12, 459-476.
- Simmons, Robert P. *Computations from the English*. Prentice-Hall, 1984.
- von Stechow, Arnim. "Topic, Focus and Local Relevance". In W. Klein and W. Levelt (eds.), *Crossing the Boundaries in Linguistics*, 1981, 95-130.
- Thompson, H. "Strategy and tactics: a model for language production". In *Papers from the 13th Regional Meeting*,

Chicago Linguistics Society, 1977.

Valain, M. "The Restricted Language Architecture of a Hybrid System", In the *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, L.A., 1985, 547-551.

APPENDIX A

In this Appendix we have given two examples of text that NOSVO can generate. The text was based upon naturally occurring text (Lawrence 1985). The old information first principle has been applied to the first text. It has not been applied to the second text. We believe that the second text is stilted, less cohesive and harder to read, though this has yet to be proven experimentally. We also believe that the missapplication of the old information first principle would be worse than its nonapplication. These are topics left for future research.

TEXT I

Long before I was tall enough to ride on the big coaster myself, I spent many pleasant hours persuading my reluctant father to accompany me.

As an aficionado of amusement parks I was overjoyed when our whole family finally flew to California to tackle Walt Disney's extravaganza.

More than two decades later, I'm still journeying to parks. (page 4)

TEXT II

I spent many pleasant hours persuading my reluctant father to accompany me long before I was tall enough to ride on the big coaster myself. I was overjoyed, as an aficionado of amusement parks, when our whole family finally flew to California to tackle Walt Disney's extravaganza.

I'm still journeying to parks more than two decades later.

