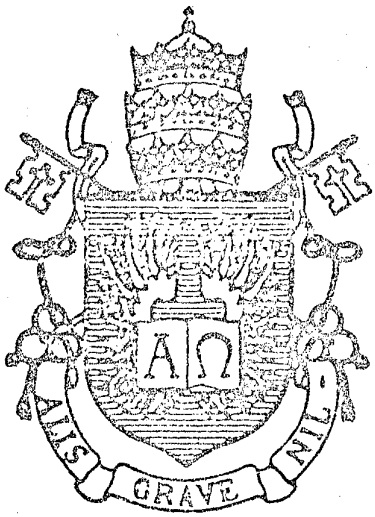

PUC

Series: Monografias em Ciência da Computação, 23/89

GETTING THE MESSAGE ACROSS: A COMPUTATIONAL GRAMMAR OF TEXT

Donia R. Scott
Clarisse S. Souza

Departamento de Informática

PONTIFÍCIA UNIVERSIDADE CATÓLICA DO RIO DE JANEIRO

RUA MARQUÊS DE SÃO VICENTE, 225 – CEP 22453

RIO DE JANEIRO – BRASIL

PUC/RJ - DEPARTAMENTO DE INFORMATICA

Series: Monografias em Ciência da Computação, 23/89

Editor: Paulo Augusto Silva Veloso October, 1989

GETTING THE MESSAGE ACROSS: A COMPUTATIONAL GRAMMAR OF TEXT

Donia R. Scott

Clarisse S. Souza

This work has been partially sponsored by FINEP

In charge of publications:

Rosane Teles Lins Castilho
Assessoria de Biblioteca, Documentação e Informação
PUC RIO, Departamento de Informática
Rua Marquês de São Vicente, 225 - Gávea
22453 - Rio de Janeiro, RJ
BRASIL

Tel.: (021) 529-9386 TELEX: 31078 FAX: (021) 274-4546
BITNET: userrtlc@lncc.bitnet

Getting the Message Across: A Computational Grammar of Text.

Donia R. Scott

Philips Research Laboratories,
Cross Oak Lane,
Redhill,
Surrey RH1 5HA,
U.K.
email: scott@prl.philips.co.uk

Clarisse Sieckenius de Souza

Departamento de Informática, PUC/RJ,
Rua Marquês de São Vicente 225,
22453 Rio de Janeiro, R.J.,
Brazil.
email: usercsds@lncc.bitnet

ABSTRACT

For text to be correct it must be grammatical, but for text to be good it must also be effective. In other words, the message must be transparent in the text; the more transparent it is, the easier it will be for the reader to understand. In order to achieve message transparency, computational models of text generation must have access to cognitively motivated text structuring rules during the planning process. These rules are related to the psycholinguistic and sociolinguistic factors that affect language processing. They belong to a text grammar, whose role is to map the desired message onto a text plan.

By proposing the existence of a text plan as an intermediary stage between message and final text, we argue for a 3-step approach to text generation and show that this brings about many theoretical and computational advantages over other approaches.

Keywords: text planning, message transparency, text grammar, style.

1. Introduction

The primary goal of text generation is that of getting the desired message across. Of course, as writers we have no way of guaranteeing that our readers have got the message. The best we can do is to try to make sure that its propositional content, and our attitude towards it, are present in the text, and that they are expressed in a way that makes them as easy as possible for the reader to decode. This latter factor is what is generally referred to as style, and is the main criterion for distinguishing good from bad writing. Given the link between style of writing and ease of reading, it is clear that the needs of the reader are of paramount importance for text generation: the more transparent the message, the better the text.

Message transparency is thus the ultimate criterion of the quality of a text. Three major factors contribute to it: the choice of discourse structure, syntax and lexis. Choices at each level must be directed by general cognitive criteria, which relate to psycholinguistic aspects of text processing. They can also be directed by more local criteria, such as register (e.g. formal vs. informal writing, letters vs. legal documents) or the needs of a particular reader (e.g. her idiolect), some of which may well override considerations of processability.

Although the need for message transparency is generally recognised in the field of automatic text generation, proposed methods for achieving it have failed to provide an adequate computational model. Plausible theoretical models have led to undesirable computational features, such as those that arise from having an interpreter within the generator (e.g. Yazdani, 1987), and computationally attractive ones have been without the theoretical basis necessary for a general computational model (e.g. Hovy, 1989).

At present, computationally tractable approaches to the transparency problem make the necessary cognitive decisions either during the planning of the message content or in conjunction with grammatical decisions. These two approaches represent opposing views of text generation. In the first, generation is a two-stage process comprising text planning and realising (McKeown, 1985; McDonald, 1984b). The output of the first stage is a full specification of the message: its propositions and their required order as defined by a precompiled discourse structure. This forms the input to the second stage,

where the grammatical rules of the target language are applied to it to produce the final text. An important point to be made here is that the first stage is completely free of linguistic knowledge and that the second has only linguistic knowledge available to it.[1] As a result, cognitively motivated linguistic structures at the discourse level are not available for direct manipulation during planning. The second approach views generation as a one-stage process, with planning and realising occurring concurrently (Appelt, 1985). Since discourse building is not an isolatable activity in the system, the discourse structure is once again not available for direct manipulation.

The inability to manipulate the discourse structure of the text directly has important consequences for the level of transparency that can be achieved. Perhaps the most important of these is that it becomes extremely difficult to produce text that is sensitive to cognitive constraints, because the appropriate decisions are either taken at the wrong time or not at all. At best, the text will be expensive and time consuming to produce. At worst, it will be of bad quality.

The main impact of not being able to mould the message into a cognitively desirable form is that the message and text plan become one and the same, and thus each message will have its own unique expression — one that is unaffected by contextual factors. Moreover, since the structure of the message will necessarily reflect the machine's reasoning processes, so too will the text. There is no reason to believe that access to a machine's "flow of thought" in constructing a message will be at all helpful for understanding it, and many reasons to believe that it will not. Finally, without access to psycholinguistically motivated constraints on syntactic operations such as embedding and coordination, decisions regarding sentence scoping can only be based on *ad hoc* rules or on strictly grammatical ones.

Making the necessary cognitive decisions, but at the wrong time, means that although one does have the opportunity to produce good text, one will have to pay a high computational price for it. All untimely decisions will have to be undone, and "undoings" are likely to have combinatorial effects. The longer the text, the higher the price is likely to be. The extreme case of making decisions too early would be a generator of the type suggested by Yazdani (1987) which produces text, then reviews it by analysing it and re-doing bits (or, in the worst case, the entire text).

---

[1] To the extent that it is generally referred to as the "linguistic component" (see e.g. McDonald, 1984b).

For opportune decision-taking to occur, the generator has to be able to recognise the points at which cognitive and linguistic knowledge contribute individually to its task. It will also have to make use of an architecture that reflects the independence of levels of processing: message, text structure and surface text. This leads to a three-stage view of text generation, shown in Figure 1: (1) deriving the message, (2) deriving the text structure to be imposed on it, and (3) writing it. Mapping the second stage onto the third demands an adequate generative grammar of the target language, of the type captured in NIGEL (Mann and Matthiessen, 1985), and MUMBLE (McDonald, 1984a). Mapping the first onto the second demands a computational grammar of text that captures the psycholinguistic and sociolinguistic factors affecting human text processing.

2. A Computational Grammar of Text

The term *text grammar* is often loosely used to refer to the grammar of the target language that underlies the process of text realisation at sentential level, such as those mentioned above. Its more specific use relates to a higher level of text processing, of the type required to map the message onto a discourse structure, where psychological and sociological aspects of language are treated (see e.g. Van Dijk, 1972; Rumelhart, 1975; Levy, 1979). Our concern in this paper is with the second of these. For us, a text grammar is a set of transformational rules which apply to the message and which result in a text plan. This text plan is a syntactically annotated discourse structure that is hierarchical in form, with the propositional content of the message as terminal elements.

The text grammar is used to direct the production of the most suitable text plan. The criterion of suitability is twofold: conforming to conventional form and exploiting the characteristics of human cognition. An example of the first would be the requirements for producing a formal wedding invitation; an example of the second would be the psycholinguistic finding that parallel structures are more easily processed than asymmetric ones (Frasier et al., 1984). Both criteria apply to all types of texts but with different degrees of relative importance. For example, in the case of legal documents, it is clear that considerations of register often override psycholinguistic ones.[2]

[2] For this reason, theoretical proposals of text coherence that are belief-driven are unable to provide a satisfactory account of documents of this type (Mann and Thompson, 1987b).

The effect of a text grammar on the generated text will be seen at all levels. Some rules will direct the structure of the discourse, and others will direct syntactic or lexical choices. Furthermore, the impact of a rule that applies to one level may be felt at another. The result of the application of the grammar to the message will be a plan that partially commands how the message must be expressed by the realisation component. Since the text plan is not sufficiently complete for direct translation into text, the realisation component must reason about how to express some of its elements. Whereas reasoning for the production of the text plan is based on information about the communicative setting in which the text will occur, reasoning performed by the realisation component is based on strictly linguistic information. Realisation rules capture the grammar of the target language and linguistically motivated stylistic rules — rules such as "avoid unnecessary repetition of words and marked syntactic structures".[3] It follows from this characterisation that text grammars are more closely aligned to theories of language processing than to issues related to particular languages.

Since the text grammar is a set of meaning-preserving transformations applied to the message, all aspects of the message must be available for treatment: its propositional content, the rhetorical relations that hold between propositions and between parts of the message, and the relative importance of the elements of the message *at all levels*: within and between both propositions and rhetorical relations. That is to say, the message itself should have a fully specified discourse structure, of the type proposed by Grosz and Sidner (1986). The reason for this is easy to see: the more cohesive and coherent the message, the easier it will be to make a sensible text out of it.

Among the many problems that arise when attempts are made to generate text from an impoverished message, are those of introducing focus and of sentence scoping. With regard to focus, the integration of independent discourse and focus structures makes it difficult to achieve context sensitivity. An example of this can be found in Hovy and McCoy (1989). The problem of sentence scoping is even more difficult to resolve in a principled way. As stated by Hovy (1988), it becomes a major problem, one that can only be made tractable by applying *ad hoc* solutions.

---

[3] See Scott and Souza, 1989, for a more extensive discussion of this.

2.1. Current Text Planners as the Basis for a Text Grammar

At present, the most fruitful approaches to text planning are the schema-based approach developed by McKeown (1985) and Rhetorical Structure Theory (RST) developed by Mann and Thompson (1985, 1987a,b). Both approaches view coherence as the maintenance of a hierarchical structure bound together by rhetorical predicates. As such, they provide a useful foundation on which to consider building a text grammar.

Schemas are skeletal discourse structures which are used to guide the selection of the propositional content of the message. They represent text patterns that are typical for a given domain and provide static discourse patterns for generation. Schemas are an attractive option for generation since they guarantee coherence and, since they are precompiled, do not introduce high computational overheads. These benefits, however, are counterbalanced by a number of limitations on the quality of text. Precisely because they are static, the freedom with which context sensitive effects can be incorporated during planning is severely curtailed. The result of this is text that is correct but which may not be easy to understand. One way around this problem would be to have access to the full set of schemas that would apply to the set of all possible contexts. However, the effect of this would be to negate the major attraction of schemas — their computability. Furthermore, since the process of schema construction is to date a manual one, and one that is extremely difficult and time-consuming, achieving the full set of schemas would be an impossible task. Automating the process, thereby making the task feasible, would require a grammar of text, the availability of which would remove the need for schemas in the first place.

RST, on the other hand, has a number of attractive features as the basis of a text grammar. First and foremost, it provides the necessary flexibility for the dynamic construction of novel text plans since its structures are compositional. Second, it has the advantage of being belief driven, thus making rhetorical relations amenable to be reasoned about. RST also allows for the possibility of linking the actions of the writer with the perceptions of the reader; rhetorical relations are defined in terms of the propositional content of its elements, the speaker's perlocutionary goal and the expected effect on the reader. In addition, RST incorporates a means for signalling the relative importance of the various parts

of the message. But RST is not without disadvantages. The most important of these is that, whilst it provides a reasonable descriptive analysis of text, it has not yet been sufficiently formalised to be used for generation within the framework proposed in this paper. It is also incomplete in that it lacks much of the linguistic knowledge required for good text generation. In the remainder of this paper, we describe some of the steps we have taken in an attempt to fill this gap.

3. A Text Grammar Based on Rhetorical Structure Theory

For RST to be useful as a generative grammar of text, adjustments are required to account for two major phenomena: 1) the interactions between the syntax and semantics of rhetorical relations, and 2) the equivalence of different configurations of rhetorical relations. In the case of the first, a full specification of the morphosyntactic constraints on relations and their elements is required. This will affect many of the decisions that are manifest at the sentence level, particularly those to do with clause combining. In the case of the second, what is required is a specification of equivalences among rhetorical structures, whose correspondences are reflected in the text at the level of discourse. Since the equivalence rules operate on complex hierarchical structures, they must be transformational rules applicable to RST trees.

Although some projection of syntactic form can be made from the specifications of rhetorical relations as provided by Mann and his colleagues (see, e.g. Mann & Thompson, 1987a, Matthiessen & Thompson, 1987), these are not sufficient for generative purposes since they do not account for many of the constraints that apply to relations and their elements. Among them are the restrictions on the syntactic constituency of nucleus and satellite, the order in which these elements may occur, the morphosyntactic markers of a relation, and the interactions between all three. Without explicit statement of such constraints, the expression of relations becomes problematic in that the derived grammar will overgenerate. The original proposal thus predicts what *can* be said, but not what *cannot* or *should not* be said. This is best illustrated with an example.

Consider the propositions:

(1) My car isn't British.

(2) My car is a Renault.

placed in an evidence relation with (1) as nucleus and (2) as satellite. The unextended theory could produce the following appropriate utterances:

My car isn't British. It's a Renault.

My car is a Renault. It's not British.

Since my car is a Renault, it's not British.

My car is not British since it's a Renault.

Because my car is a Renault, it's not British.

My car is not British because it's a Renault.

My car is a Renault, therefore it's not British.

It will, however, also produce a large number of inappropriate ones, such as:

*Since my car is not British. It's a Renault.

*Therefore my car is a Renault. Since it's not British.

*My car is not British, therefore it's a Renault.

*Because my car is a Renault, since it isn't British.

*Since my car isn't British, because it's a Renault.

and the like.

It is clear that overgeneration is to be prevented. One way of doing this would be to include the above mentioned clause combining constraints in the specifications of relations, thereby blocking the possibility of deriving inappropriate structures. Another would be to filter out the spurious derivations at the realisation stage. This latter option would, however, entail theoretical and computational difficulties that are not presented by the former. In theoretical terms, it would involve a mixing of levels since decisions regarding the best syntactic form must depend on their suprasentential consequences — information that is not available at the realisation stage. In computational terms, structures will be built, only to be later abandoned. The related processing and storage costs will

increase exponentially with the size of the text to be generated. For all these reasons, we have chosen to opt for the first possibility.

The ability to generate different texts to convey the same message presumes the existence of equivalence criteria. The second requirement in an expanded RST text grammar is therefore that of capturing structural equivalences among text plans. This amounts to specifying the set of meaning preserving transformational rules that produce two kinds of structural variations: those which maintain the rhetorical relations of the message but not their configuration, and those which maintain neither. Examples of these are shown in Figures 2 and 3 respectively. The heavy lines in the figures correspond to the nuclei of relations, and the light lines to satellites.

The text grammar controls the transformations that can be made to the structure of the message to produce a text plan, and commands the types of syntactic structures that must be produced by the realisational component. These syntactic decisions are made by reference to the syntactic specifications found in the definition of the relations and by the operation of structure-sensitive, meaning preserving syntactic rules. The text plan then is an RST structure that is semantically equivalent to the message and whose terminal elements are syntactically annotated.

The grammar operates by traversing the message structure in a bottom-up manner starting with the most nuclear element of the message. By consulting the syntactic specifications of the elements of each relation it finds, and the context of the current state of the text plan, it determines the order of presentation of the elements of a relation, and their appropriate syntactic categories. In addition, it determines the the appropriate location for attachment to other relations. These decisions are affected by a number of factors. For example, the chosen order of the elements of a relation is also affected by their complexity (ie. the number of other relations they contain), and the strength of the relation *per se* (to be discussed in more detail below). Also, the syntactic annotations of terminal elements are affected by pragmatic aspects of the content of the terminal elements.

The grammar we have developed places constraints on the way messages can be structured and on the way their propositions should be represented. For example, the message planner should attach weak relations at a higher level in the message structure than strong ones (the reason for this will be

made clear below) and domain knowledge must be respresented in a way that makes it possible to produce a possible sub-clausal realisation.

3.1. Some text grammar rules

In what follows, we present a set of equivalence rules of the type required to produce the sorts of transformations between message and text plan shown in the figures. The rules we give do not comprise the full set of possible transformations. Instead, they are suggestive of the types of transformations that must be accounted for, and relate only to cognitively motivated transformations. They are derived from two sources: psycholinguistic evidence on text processing and our intuitions as writers. The rules we present here form part of a text grammar we are developing with a view to generating text that optimises the chances of achieving message transparency. Given that our grammar is RST-based, it follows that the message on which it operates is itself an RST structure and that the operands of the grammar are rhetorical relations. For purposes of simplicity, we will assume here that each terminal node is a clausal proposition.

RULE 1: Only satellites can be embedded.

This rule guarantees the maintenance of the intégrity of the message throughout the transformational cycle by preventing operations that will result in the loss of the original rhetorical relations and of the salience relations between their elements. One of the effects of this rule is that embedding does not result in the decomposition of complex satellites.

RULE 2: Embedding should occur within the left-most clause in the text plan that bears a semantic relation at the clausal level to the candidate clause.

Since the nucleus of a relation may be complex, some direction must be given for the appropriate locus of embedding. This rule ensures that embedding does not lead to stylistic blunders such as:

Since Owen has an American passport, he, who was born in Jamaica, is an American citizen.

instead of

Since Owen, who was born in Jamaica, has an American passport, he's an American citizen.

Notice that this rule gives preference to the order of presentation of the elements of a relation in the text plan over their salience relations.

RULE 3: Satellites of "weak" relations should, where possible, be embedded, provided that the number of remaining clauses is not 1.

As we stated before, the relevance of its component propositions is an important part of the message, and this information must obviously be signalled. Rule 3 ensures that a "weak" part of the message is not made prominent in the text. A useful indicator of the strength of a relation is via its morphosyntactic markers. For example, Elaboration and Circumstance are among the weakest rhetorical links in that the content of their satellite tends to be less central to the message than, for example, Concession and Antithesis. Elaboration and Circumstance have no markers, whereas Concession and Antithesis demand the use of a marker.

Since embedding reflects the semantic subordination of satellite to nucleus, it is therefore a particularly desirable expression for the satellite of weak relations. Other expressions of it tend to have the effect of giving more weight to the satellite than it rightly deserves, even those that include expressions such as "by the way" or "incidentally". Thus, Rule 3 commands the embedding of satellites of a weak relation. The proviso to the rule ensures that, where possible, the text does not contain dangling sentences. Dangling sentences can be severely disruptive, giving the impression of information being included as an afterthought. They can also be extremely misleading, since by trying to make sense of its isolated appearance, the reader often perceives a completely erroneous meaning — that is a new topic. The weaker the relation, and the larger its nucleus, the more pronunced the effect will be.

RULE 4: In cases where the specification of a relation permits more than one syntactic expression of its satellite, the order of preference is adjective, adverb, appositive noun phrase, prepositional phrase and subordinate clause.

Here we prescribe the preferred surface expressions of semantic embedding. Rule 4 reflects a preference for simple over complex syntactic structures, and as such relates to the Gricean maxim of quantity (Grice, 1975). Of course, the range of choice between the possible expressions of a clause will be directed by the target language. For example, embedding

Paula was in the mood to dance with Peter.

into

Paula danced with Peter.

would be realised as an adverb in English (Paula danced with Peter *willingly*) but as a prepositional phrase in Portuguese (Paula dancou com Pedro *com vontade*).

Rules 1 to 4 relate to the maintenance of the semantic subordination of satellite to nucleus during the construction of a text plan. The expression of non-subordinate structures (ie. multi-nuclear or multi-satellite ones) is dealt with in rules 5 and 6.

RULE 5: Coordination can only occur between elements of List, Sequence and Contrast.

Although coordination can, in principle, apply to the full set of multi-nuclear and multi-satellite relations, we have chosen to restrict its application to situations that guarantee the syntactic expression of semantic non-subordination. We assume that in the remaining situations, morphosyntactic markers of the relation are more appropriate.

RULE 6: The greater the number of shared parameter values between clauses, the more desirable it is to coordinate them.

Here we refer to shared features of the elements of a message at the propositional or rhetorical level — ranging from syntactic constituents, to focus elements, to relations themselves. The effect of this rule is to reinforce the general principle of conciseness discussed with respect to Rule 4. It also results in the preferential production of shorter over longer coordinated sentences. For example, it gives priority to coordinations that lead to ellipsis over those which do not.

The effect of Rules 1 to 6 on the message is to transform its structure, maintaining the given rhetorical relations, and to add syntactic annotations to the terminal elements of the structure (see

figure 2). In addition to them, we propose the existence of rules for metamorphic structural transformations of the type shown in Figure 3. Such rules are, however, extremely difficult to discover "by hand" since they can only be revealed in situations where it can be ensured that the reader (in this case the rule-writer) fully gets the message of the text being transformed. Deriving them by more formal methods would obviously be preferred. However, this would require models of belief construction of the type that present some of the greatest challenges to current research in Artificial Intelligence.

3.2. The impact of the rules on the text

Producing text without a text grammar would amount to slavishly following the structure of the message — in fact, to treating the message itself as a text plan. As an example of the undesirability of this, consider:

George received a letter from Peter and it arrived yesterday. However, he had specifically told Peter never to contact him. Peter is George's brother-in-law and an ex-con.

which is the best that one could expect for the message in figure 2. Although this is perfectly acceptable as a grammatical English text, it is certainly not the only possible textual rendition of the message. Neither is it necessarily the best. The appropriateness of a text is clearly dependent on the context in which it occurs, and although one can imagine a context where this text would be entirely appropriate, it is clear that taken in the absence of a surrounding context, the text in Figure 2 is vastly superior. By successively unfolding the message in packets of information instead of distributing it throughout the text, "re-packaging" is made much easier. The treatment of information related to Peter is an example of this. In this case, such "packaging" is the result of applying rules 3 and 4.

Other attempts to improve on the quality of text produced with a message-as-text-plan approach have treated the problem as essentially one of sentence scoping. Although this type of treatment can lead to improvements (e.g. Hovy & McCoy, 1989), a number of other problems remain unresolved. For example, without rules to direct the location and form of embedding, texts such as:

Although George had specifically told Peter never to contact him, he received a letter from

Peter, his ex-con brother-in-law, yesterday.

or

George received a letter from his brother-in-law Peter yesterday even though he had specifically told Peter, an ex-con, never to contact him.

or even

Although George had specifically told his brother-in-law Peter never to contact him, he received a letter from him yesterday. Peter is an ex-con.

could be generated from the message in Figure 2. Indeed, one of the advantages of the approach we have chosen is that sentence scoping is no longer an issue that demands a separate treatment. Instead, decisions about sentence scope follow naturally from other, higher-level, decisions.

Treating syntax simply as an expression of semantic relations provides some additional advantages. With regard to coordination, one of the advantages is that utterances such as:

**The hunter shot the breeze, the bolt and the wild boar.*

or

**The carpenter mended the table with force, glue and three nails.*

are prevented at a very early stage. Another advantage is that knowledge about those situations which appear to be cases of coordination, but which are not, is available. The generator can therefore reason about situations where the significance of an ambiguous marker of discourse and syntax may be misunderstood. For example, the word *and* in:

They're not getting along well and she's had to move out.

is a lexical marker of a causal relation in:

I saw Anna and Peter yesterday. *They're not getting along well and she's had to move out.*

but a marker of syntactic coordination in:

Poor Anna! She's having problems both with her relationship with Peter and with her

apartment. *They're not getting along well and she's had to move out.*

Metamorphic rules, of the type suggested in Figure 3, play an important role in the maintenance of the thread of discourse. Their use is motivated by notions concerning the suitability of a particular configuration of rhetorical relations to its preceding context. This can be seen in the following examples, where the italicized portions of the text convey the same message but with structures more suitable to the context.

> Young stockbrokers do well for themselves these days. *Simon is a young stockbroker. In other words, he's a yuppie. He works in Central London, lives in Chelsea, has a filofax, and drives a Porsche.*

> These days, being a young stockbroker is almost synonymous with being a yuppie. *Simon is a young stockbroker. He works in Central London, lives in Chelsea, has a filofax, and drives a Porsche. In sum, he's a yuppie.*

The Restatement and Summary relations are prevalent in argumentative texts. "Making the point" is in fact the whole point of these kinds of texts, and some structural configurations tend to be more effective for a given type of argumentation than others. The wrong choice of structure will hinder the recovery of the argument and thus the scoring of the point. This would be the effect of switching around the contexts in the previous examples:

> These days, being a young stockbroker is almost synonymous with being a yuppie. *Simon is a young stockbroker. In other words, he's a yuppie. He works in Central London, lives in Chelsea, has a filofax, and drives a Porsche.*

> Young stockbrokers do well for themselves these days. *Simon is a young stockbroker. He works in Central London, lives in Chelsea, has a filofax, and drives a Porsche. In sum, he's a yuppie.*

4. Discussion

It is our view that the most appropriate theoretical framework for the computational treatment of message transparency is one which provides a clear separation between issues of *language theory* and

of *language type*. From the theory of language, we derive a text grammar that guides the production of a cognitively oriented text structure. From the grammar of the target language, we realise the text as an object for perception. Clearly, since text grammars are intended to express global cognitive principles, they should be expected to hold universally for all languages. With regard to the grammar we have described, extensive experimentation with two unrelated languages, English and Brazilian Portuguese, strongly indicates that this is the case. Although cross-language differences do occur, accounting for them amounts to fine-tuning the grammar rules and not to the production of alternative rules at the same level of abstraction.

One of the attractions of our three-stage model is that by separating issues of language perception from those of language production, we are able to handle many of the issues that affect modularity and portability in a theoretically motivated, and therefore generally applicable, way. Other, more linguistic, problems are also greatly reduced by the potential to deal with them in a theoretically well motivated manner. Among these, the problems of focus and of sentence scoping are perhaps the most important.

Another advantage of our approach is that by treating text production and understanding as intimately linked processes, we are in a better position to exploit their interconnections in a full natural language interaction. By recognising a non-linguistic status of knowledge representation and by providing cognitive rules for moulding this knowledge structure into different textual forms, we imply the need for a more-or-less symmetric architecture for text understanding systems. This would call for two mappings in the analysis process: one from text to text plan, and another from text plan to message. Approaches to text understanding which postulate meaning at the level of the text plan (ie. those which do not perform the necessary mapping onto the message), make it extremely difficult to assess the correspondences between two received renditions of the same message or between an analysed input and its subsequent regeneration. As such, their behaviour is perhaps better characterised as text recognition than as text understanding. An extreme case of this would be the use of schemas for textual analysis.

The third and most important advantage of our approach is that it leads to the production of text

that is qualitatively better than that produced by other approaches. One exception to this would be text produced by systems of the type suggested by Yazdani (1987), which could be of an equivalent standard. Their approach is, however, more appropriate to "batch-job" type applications, such as critiquing and style-checking, than to text generation in a interactive setting.

Psycholinguistic influences on text production have long been of concern to researchers in text generation. Attempts to account for them have, however, been more closely aligned to the language producer than the language consumer, that is, to language performance than language competence. Although they have led to psychologically plausible computational models of language generation, capable of reproducing many of the performance features of the human language producer (see, e.g., McDonald, 1984a,b), the text produced will often not lead to easy retrieval of its underlying message. This is hardly surprising, given that as speakers and writers we are extremely bad performers — our speech is often littered with hesitations, false starts and so-called slips of the tongue, and few of us are able to produce a perfect text in one go.[4] The more closely aligned a model of text generation is to the needs and characteristics of readers (and to a theory of language competence) the higher its potential will be for achieving message transparency.

Of course, message transparency can never be guaranteed, neither for the human nor machine writer, since readers' perceptions are inevitably affected by their own knowledge and experience. These factors are impossible to control, especially when writing for a general audience. Good writing, then, means not adding any extra "noise" to the understanding process. We do this by attempting to tailor the message of the text, and the terminology used to express it, to our intended audience. We also do it by structuring the text in a way that we think will be most effective. The first two of these amount to some sort of personalisation of the text. In terms of HCI, this is dependent on the accuracy of the system's model of the user. Achieving the third means appealing to rather more general characteristics of readers, by exploiting the cognitive principles that are presumed to be intrinsic to human text processing. In computational terms, this requires the use of a psycholinguistically-based grammar of text, such as that proposed in this paper.

[4] Fromkin (1980) and Cutler (1982) provide an excellent discussion of these phenomena.

An important feature of grammars of text is that they prioritise *communication* over *linguistics*. Assuming the perspective of the reader implies the need for a conciliation, on the part of the speaker, between decisions about what to say and our hypotheses about how it is best communicated. This entails a shift of the emphasis of text generation from that of exploiting the limits of grammaticality to that of exploiting the limits of comprehension, from criteria based on notions of text grammaticality to those based on notions of message transparency. One consequence of this is that some aspects of text which were previously thought to be peripheral now come to be seen as playing a rather more central role. For example, style is no longer a matter of aesthetics but of cognition.

ACKNOWLEDGEMENT

We are extremely grateful to Maria das Gracas Volpe Nunes for her contribution to many of the ideas discussed here.

5. References

Appelt, D. (1985). *Planning English Sentences*. Cambridge: Cambridge University Press.

Cutler, A. (1982). *Slips of the Tongue and Language Production*. Amsterdam: Mouton.

Frazier, L., Taft, L., Roeper, T., Clifton, C., Ehrlich, K. (1984). Parallel Structure: a Source of Facilitation in Sentence Comprehension, *Memory and Cognition*, 12.

Fromkin, V. (1980). *Errors in Linguistic Performance: Slips of the Tongue, Ear, Pen and Hand*. New York: Academic Press

Grice, H.P. (1975). Logic and Conversation. In P. Cole and J.L. Morgan, eds., Syntax and Semantics: Speech Acts. New York: Academic Press.

Grosz, B.J. and Sidner, C.L. (1986). Attentions, Intentions and the Structure of Discourse. *Computational Linguistics*, Volume 12, Number 3, pp 175 - 204.

Hovy, E. (1988). Planning Coherent Multisentential Text. *Proceedings of the 26th Meeting of the ACL*. Buffalo, New York.

Hovy, E. (1989). Pragmatics and Natural Language Generation. To appear in *AI Journal*. Also available as USC/Information Sciences Institute Research Report RS-89-233.

Hovy, E. and McCoy, K.F. (1989). Focusing Your RST: a Step Toward Generating Coherent Multisentential Text. In Proceedings of the 11th conference of the Cognitive Science Society, Ann Arbor, August 1989.

Levy, D.M. (1979). Communicative Goals and Strategies: Between Discourse and Syntax. In T. Givon, ed., *Syntax and Semantics: Discourse and Syntax*, New York: Academic Press.

Mann, W.C. and Thompson, S.A. (1985). Assertions from Discourse Structure. USC/Information Sciences Institute Research Report RS-85-155.

Mann, W.C. and Thompson, S.A. (1987a). Antithesis: a Study in Clause Combining and Discourse Structure. USC/Information Sciences Institute Research Report RS-87-171.

Mann, W.C. and Thompson, S.A. (1987b). Rhetorical Structure Theory: a Theory of Text Organization. In Livia Polanyi, ed., *The Structure of Discourse*, Norwood, N.J.: Ablex Publishing

Corporation.

Matthiessen, C. and Thompson, S.A. (1987). The Structure of Discourse and 'Subordination'. In John Halman and Sandra A. Thompson, eds., *Clause-Combining in Discourse and 'Subordination'*. Amsterdam: John Benjamins Publ. Co.

McDonald, D.D. (1984a). Description Directed Control: Its Implications for Natural Language Generation. In N. Cercone, ed. *Computational Linguistics*. New York: Plenum Press.

McDonald, D.D. (1984b). Natural Language Generation as a Computational Problem: an Introduction. In M. Brady and R. C. Berwick, eds., *Computational Models of Discourse*, Cambridge, Mass.: MIT Press.

McKeown, K.R. (1985). *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*. Cambridge: Cambridge University Press.

Mann, W.C. and Matthiessen, C. (1985). A Demonstration of the Nigel Text Generation Computer Program. In J. Benson and W. Greaves, eds., *Systemic Perspectives on Discourse*, Volume 1. Norwood, N.J., Ablex Publishing Corporation.

Rumelhart, D.E. (1975). Notes on a Schema for Stories. In D.G. Bobrow and A. Collins, eds., *Representation and Understanding*. New York: Academic Press.

Scott, D.R. and Souza, C.S. (1989). Conciliatory Planning for Extended Descriptive Texts. Technical Report 2822, Philips Research Laboratories, Redhill.

Yazdani, M. (1987). Reviewing as a Component of the Text Generation Process. In G. Kempen, ed., *Natural Language Generation*, Dordrecht: Martinus Nijhoff Publishers.

Van Dijk, T.A. (1972).R *Some Aspects of Text Grammars*. The Hague: Mouton.

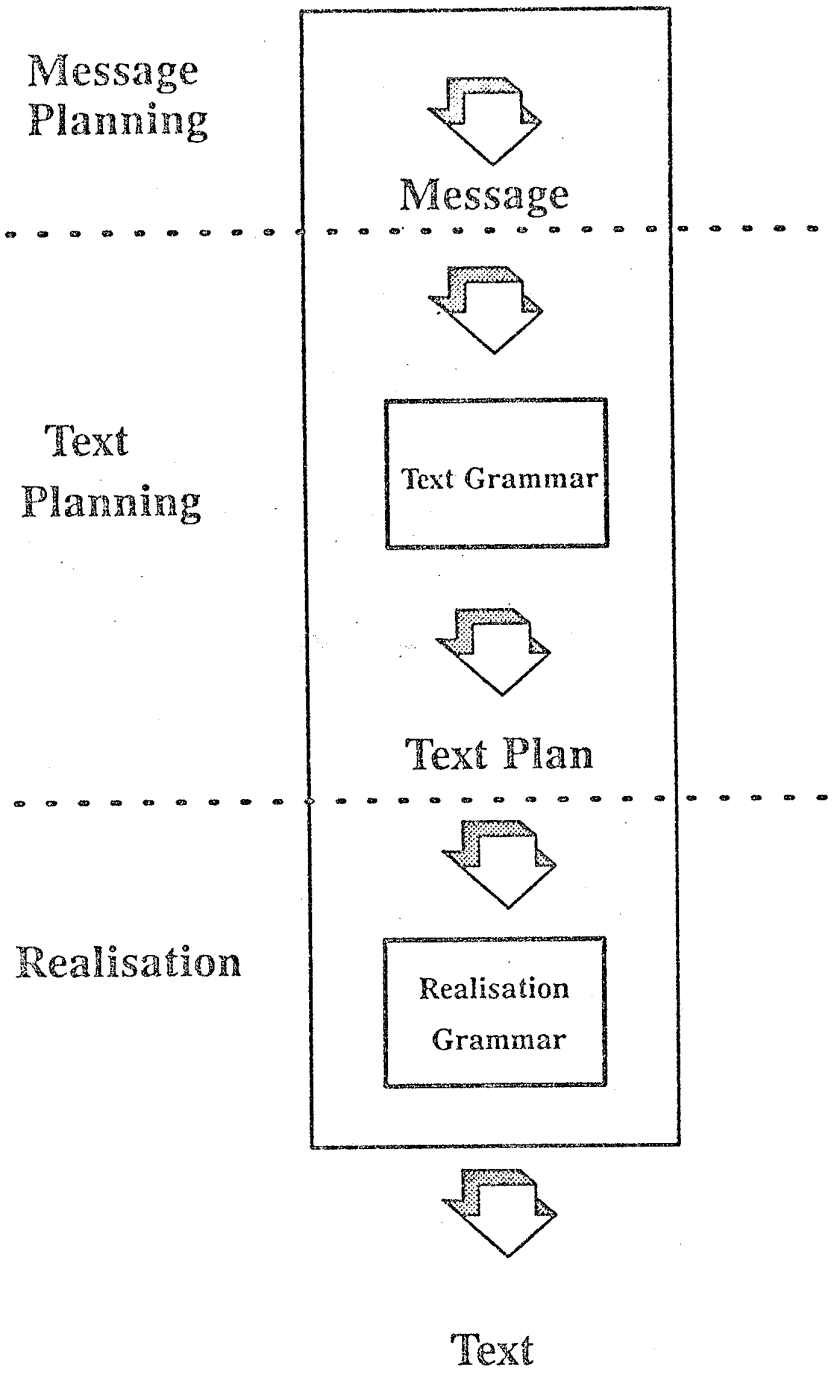

Figure 1: Conciliatory Planning via a 3–step approach.

Propositions

(a) George received a letter from Peter.

(b) The letter arrived yesterday.

(c) George had specifically told Peter never to contact him.

(d) Peter is George's brother-in-law.

(e) Peter is an ex-con.

Text Plan

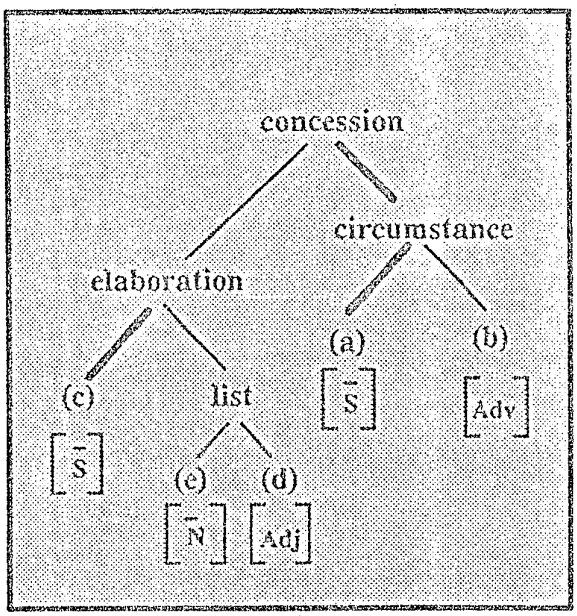

Message

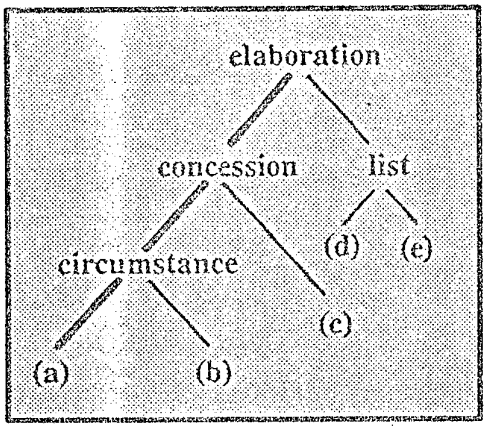

Text

Although George had specifically told his ex-con brother-in-law Peter never to contact him, he received a letter from him yesterday.

Figure 2: A transformation of message to text plan which maintains the original rhetorical relations.

Propositions:

(a) Simon is a young stockbroker.

(b) Simon is a yuppie.

(c) Simon works in Central London.

(d) Simon lives in Chelsea.

(e) Simon has a filofax.

(f) Simon drives a Porsche.

Message

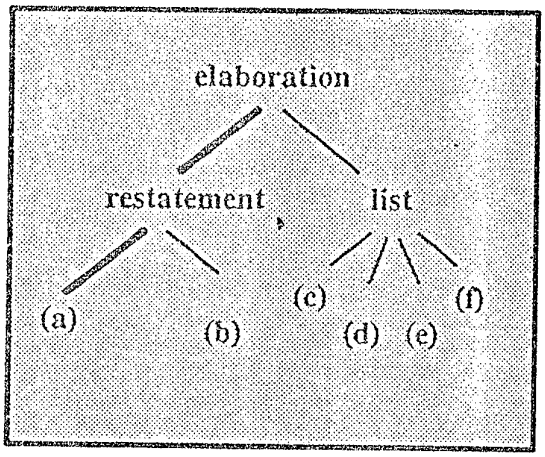

Text Plan

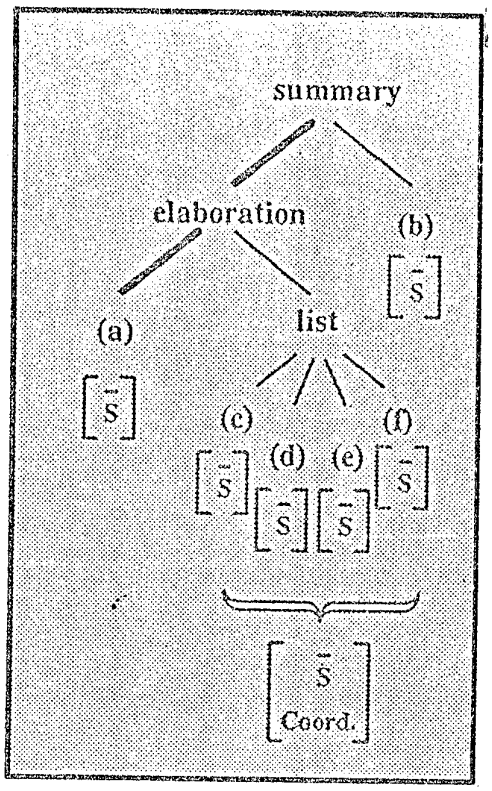

Text

Simon is a young stockbroker. He works in Central London, lives in Chelsea, has a filofax and drives a Porsche. In sum, he's a yuppie.

Figure 3: A transformation of message to text plan which does not maintain the original rhetorical relations.