ATLANTIS XIV 1-2, May-November 1992 (81-105)

INDICE

A COMPUTATIONAL APPROACH TO THE THEORY OF DISCOURSE STRUCTURE: ADVANCES IN TEXT PLANNING

Julia Lavid
Universidad Complutense

The aim of this paper is to introduce to the linguistic community some preliminary findings in the area of text planning which I hope will serve to lay the foundations of a computational theory of discourse structure. The current text planner, an ongoing research project, presents in a declarative way some of the linguistic resources necessary to generate coherent discourse, including communicative goals, text types, schemas, discourse structure relations, and theme selection patterns. This knowledge is encoded as separate resources which are activated by a flexible planning process that accesses appropriate information from its different linguistic resources and gradually constructs a discourse structure tree. The paper decribes the resources which have been successfully identified and represented to date, and the planning mechanism that activates them.*

1. Introduction

One of the main differences between the research goals of linguistics and computational linguistics is the attention given by the latter to processes. While linguistics has been traditionally oriented towards the description, analysis and explanation of linguistic phenomena, computational linguistics cannot stop there but needs to develop pro-

* The ideas presented in this paper are based on discussions of the ISI text planning group consisting of Eduard Hovy, Cecile Paris, Elisabeth Maier and Vibhu Mittal. I would like to give my special thanks to Dr. Eduard Hovy for comments and revision of an earlier draft of this paper.

cesses that use the information provided by linguistic studies in order to actively "understand" and "generate" language by computer. This dichotomy of interests has given rise to general debates in computer science regarding the so-called issues of "declarativity versus procedurality", that is, the split between the static, declarative representation of knowledge and the processes that act on that static knowledge.

One of the areas of concern in Computational Linguistics is the generation of coherent multisentential text from computer-internal data formats. The process of assembling material and organising it into a coherent whole is called *text planning*. Text planners use declarative knowledge to guide the planning process in the construction of a tree-like structure that represents the discourse organization.The current understanding of the text planning process is one of activation of the declarative resources available to a text planner; the process takes decisions based on those resources. Ideally, all the declarative resources should embody all the linguistic information necessary to produce a text. Many sentence generation systems (*Penman* 1989; Hovy 1990) represent their resources in a declarative form: the computational systemic grammar *Nigel* can be considered as a declarative resource. Likewise, semantic information used in sentence generation is usually represented declaratively in the so-called knowledge base. Unfortunately, in many current text planning systems, much of the linguistic knowledge is not represented explicitly, but is still implicitly encoded in the planner itself, resulting in complex and opaque systems, as well as systems that eventually become harder to maintain and augment. Therefore, it is necessary to clearly separate linguistic knowledge from procedural knowledge in a text planning system, with the following goals in mind:

- Enhancing modularity and transparency: when linguistic knowledge is explicitly represented, it is easier to examine the resources independently of the processes using them.
- Increasing extensibility and modifiability: because of the modularity and declarativity, it is easy to isolate all

knowledge related to some specific phenomenon (such as, say, theme), to make changes and to extend the system.

- Increasing generalizability: when knowledge is explicitly represented in a clear fashion, generalizations are easier to form and incorporate.

2. Linguistic Knowledge and the Control Mechanism

What do we mean when we talk about linguistic knowledge in text generation? Generally speaking, linguistic knowledge refers to all the resources available to the text planning system when constructing a text. These resources represent the linguistic potential made available by the language, such as rules governing coherence, theme development, cohesion, inclusion of appropriate material, etc. At the text planning level, these resources cover different types of information: text types, communicative goals, discourse relations, textual statuses (theme development, newsworthiness) etc., as we will see below.

It is not enough, however, simply to build a collection of rules that embody linguistic knowledge. In order to apply them a text planning system needs to act on the linguistic information represented statically in a declarative way: it has to take decisions and choose the alternative preferred; other times it has to assemble information from different resources and present a specification of the interaction among them. This process is called in computational linguistics the "control mechanism".

In some computational work at the level of single-sentence generation there has been an attempt towards separating the grammar and the lexicon (represented declaratively as static resources) from the actual generation process (Emele et al. 1990; Emele et al. 1991). However, at the level of text organization this endeavour has not been so successful to date. A reason for this is that not all the linguistic information has been clearly identified and represented in an explicit and modular way; it is still rather unclear what the resources and their interactions are and how one can represent them declaratively. Let's concentrate on the area of discourse and review what type of linguistic information has been used in current text planning systems.

3. Linguistic Knowledge Captured in Recent Text Planning Systems

One of the major challenges for any text planning system has always been how to achieve coherence. Though most generators were capable of simply stringing together two or more sentences, they did not contain mechanisms responsible for guaranteeing a coherent paragraph; coherence was achieved as a lucky by-product of the structure of the data presented. When the data was unstructured, incoherent text resulted. As a consequence, computational researchers investigated the issue of coherence from two different perspectives:

a) the "discourse structure" approach to coherence;
b) the "content" approach to coherence.

a) By the "discourse structure" perspective we mean those attempts which have tried to give an account of coherence in terms of the relations existing between the successive pieces of a text. In the literature, these relations have variably been called rhetorical predicates (Grimes 1975; Meyer 1975), rhetorical relations (Grosz and Sidner 1986; Mann and Thompson 1987), coherence relations (Hobbs 1979). In the Artificial Intelligence arena, Hobbs (1978; 1979; 1982) was one of the pioneers in postulating a set of relations organized into four categories, which he claimed to be the four types of phenomena that occur during conversation. However, his categorization was incomplete and not well-motivated. McKeown (1982) took a slightly different approach in her thesis work, when she defined a set of static schemas that represent the structure of stereotypical paragraphs for describing objects. Coherence was enforced by the correct nesting and filling-in of these schemas, but no explicit theory of coherence was offered.

Mann and Thompson (1983; 1987) developed a descriptive theory of the hierarchical organization of natural texts which they called Rhetorical Structure Theory (RST). The theory characterized text structure in terms of approximately 25 rhetorical relations that normally occur between portions of a coherent English text, such as Purpose, Elaboration, Cause. Their claim was that a text is coherent only if all its parts can be made to fit into one overarching relation. The relations proposed include most of Hobbs's relations and support McKeown's schemas, but the theory was purely descriptive and no formal definition of the relations was given.

As defined by its authors, each RST relation holds between two non-overlapping text spans: the nucleus (N) and the satellite (S). The nucleus is that item in the pair that is most essential to the writer's purpose; the satellite presents the information that supports, elaborates, etc., in a word, that is subsidiary or ancillary to the material presented in the nucleus. Each part has a set of constraints on the entities that can be related. Relations may also have requirements on the combination of the two parts. In addition, each relation has an effect field, which is intended to denote the conditions which the speaker is attempting to achieve.

Hovy (1988a; 1988b) showed that by reformulating the relations as plan operators, each of which has a specific goal, RST can actually be used for computationally generating not only paragraphs of written texts, but also open-ended "conversation-like behaviour". This is achieved by including into RST relations so-called growth points, which act as suggestions for expansion of the tree. Hovy called his system a "text structure planner"; it was linked to various host systems (such as a data base question-answering system) which supplied it with a collection of facts to communicate and one or more communicative goals. The text structure planner constructed one or more paragragh trees and submitted the leaves of this tree, each a clause-sized piece of material, to the sentence generator PENMAN to be generated.

Later on, Moore (1989) and Paris (1988) extended and refined Hovy's structurer by suggesting that the same planning method could be used simultaneously to collect the material to be generated from an expert system. They developed a sophisticated plan language and an extensive library of plans, and used them in their text planner, a system which has influenced much subsequent work (Cawsey 1990; Maybury 1990).

b) The second approach to coherence focuses on the *content* of the discourse segments. In linguistics, this type of coherence has been called referential or topic continuity (cf. Garnham, Oakhill and Johnson-Laird 1982; Givon 1983). In this approach a discourse is coherent if there is repeated reference to the same set of entities, for instance via argument-overlap (Kintsch and van Dijk 1978), if there is a certain semantic congruence between two discourse units (Polanyi 1986), or if there is a pattern corresponding to stereotypical situations, such as visiting a restaurant or a birthday party (Schank and Abelson 1977). The notion of *focus* in computational linguistics is the Artificial Intelligence attempt to capture this phenomenon. *Focus* has been traditionally considered in Artificial Intelligence as that element on which the participants center their attention as the discourse unfolds. This element is the focus of discourse and the active process on the part of the speaker and the listener by which they concentrate their attention on some subset of their knowledge is called focusing. For example, in a sentence like "I got a really pretty turtle this weekend", the focus is the turtle, because it is more likely that the participants will talk about it in the next sentence (Sidner 1983, 116). In general, most of the work has emphasized the use of focus as a working concept for tracking discourse participants in conversation, i.e, for the resolution of anaphora, thus for aiding the understanding of a text for a particular task (Hobbs 1979; Linde 1979; Reichman 1984; Carberry 1983; Reichman-Adar 1984; Grosz and Sidner 1985; Cohen 1987).

Both the text structuring attempts and the studies on focusing phenomena have led to very useful research both in linguistic discourse studies (Redecker 1990; Sanders and Spooren 1991; etc.) and in computational text generation. But there are serious shortcomings. Text structuring is one of the pre-generation planning tasks, but not the only one. Likewise, focus is still an intuitive attempt to capture a range of textual phenomena (theme-rheme, given-new information) that text generation systems have not dared to address yet. The new text planner we have been developing at ISI (Information Sciences Institute of the University of Southern California in Los Angeles) aims precisely at bridging this gap, putting some serious linguistic knowledge into a computational framework, building upon lessons learned from predecessors in the area.

4. Towards the Future: The New Text Planner

4.1. *Identifying some of the linguistic resources needed*

The new text planner is an ongoing project at ISI and at IPSI (in Darmstadt, Germany) and it aims to lay the foundations of a theory of discourse structure by identifying and representing in a declarative way the linguistic resources, on the one hand, and by providing an adequate control mechanism that uses and activates them, on the other.

The idea is that the text planner (which is an active process) appeals to different types of resources when having to construct the meaning and organization of a text. The linguistic resources are like an imaginary map that the planner, a dynamic process, has to read when generating a text, by taking decisions at given points. Using information gathered from different sources, the planner follows a route on the imaginary map, disregarding others. That is the reason why the more explicit and motivated the resources are the easier it is for the planner to choose alternatives at each decision point.

Given the complexity of the information needed to organize a text, and the plethora of types of information available, building a collection of information resources is a daunting problem. Fortunately, the equivalent problem at the grammatical level has been studied. One way to organize linguistic resources at the grammatical level is the metafunctional split proposed by

Halliday: according to him, three types of meaning occur in language: ideational, interpersonal and textual meaning. This threefold characterization can be used as motivation for the development of an architecture for text generation systems as described by Matthiessen and Bateman (1991) and Matthiessen (1987). These authors distinguish three bases or modules where the three different kinds of meanings can be integrated: the ideation base which concentrates on experiential knowledge, the interaction base which studies interpersonal phenomena, and the text base that accounts for such textual resources as theme-rheme structures, newsworthiness,etc.

The ideation base has already been studied in detail: for example the so called "Upper Model" of the text generation system PENMAN (Mann 1983) is a linguistically motivated taxonomy of the world to which different applications can be subordinated by means of a domain specific knowledge, the "Domain Knowledge".

Both the text base and the interaction base are still under development.[1] The text planner, which has the function of organizing the text, falls into the text base. Its role in this threefold architecture is to incorporate and use resources that appeal to these three different meanings, and to integrate them during the text planning process. We

[1] The integration of interpersonal or pragmatic factors in text planning systems is still in an infant state, although some attempts have been made in the area: Cohen (1978) studied the effect of the hearer's knowledge on the selection of appropriate speech acts; Appelt (1985) described planning the inclusion and organization of material in sentences based on hearer knowledge; the effect of hearer knowledge on object descriptions was studied by Paris (1987) and on user instruction by Woolf and McDonald (1984); the explanation generator of Swartout (1983) had a switch distinguishing between two types of hearer knowledge (either programmer or medical expert). Jameson (1987) describes a program that selects appropriate utterances in evaluative contexts such as job interview situations, based on what effect each utterance is defined to have on the hearer's belief state. The program ERMA (Clippinger 1974) contains a module that reasons about the pragmatic effects of its generation goals. The most complete attempt to introduce pragmatic constraints in a generation system was PAULINE (Hovy 1988c). The system contained approximately a dozen of so-called Rhetorical Goals, including partiality, force, haste, detail, verbosity, as well as strategies that control text generation decisions to achieve the appropriate stylistic effects. We hope to take advantage of some of these findings to integrate them as another resource in our text planner.

start with some of the resources that we have identified, knowing that there are still many left that we will be incorporating in the future. To date, we have identified the following:

1.1.1. A taxonomy of discourse structure relations

The discourse structure relations are extensions of the RST relations and others, collected and taxonomized by Maier and Hovy (1991) and Hovy (1990), and merged with the conjunctive relations of Martin (1991).

At the top level, this taxonomy splits into three groups following the metafunctional breakdown of resources mentioned before. The motivation behind it is that these relations seem to link different types of meanings: *ideational, interpersonal* and *textual* meaning. Ideational relations describe how two text units are linked in terms of experiential meaning, i.e., how two text units have to be interpreted semantically with respect to each other. From a computational point of view, these relations can be defined with *respect to an ideation base, involving common sense and domain* knowledge.

Interpersonal relations are links between two textual units which address specific features about the writer and/or the reader. For example, Motivation is a relation where one text span provides arguments in order to make the reader plan or do what is presented in the other; Evaluation gives in one textual unit a writer-centered perspective for a neutrally stated fact in the other.

Textual relations serve to organize the text itself rather than the world the text describes. They don't link two text spans experientially but "rhetorically". In written monologues their function across genres is to scaffold the schematic structure of a text; in dialogues they mark challenges to moves and links between exchanges.

At present, the network of discourse structure relations contains approximately 200 relations, defined at varying levels of precision. Figure 1 shows the discourse structure relations network.

Figure -1- DISCOURSE STRUCTURE RELATIONS NETWORK

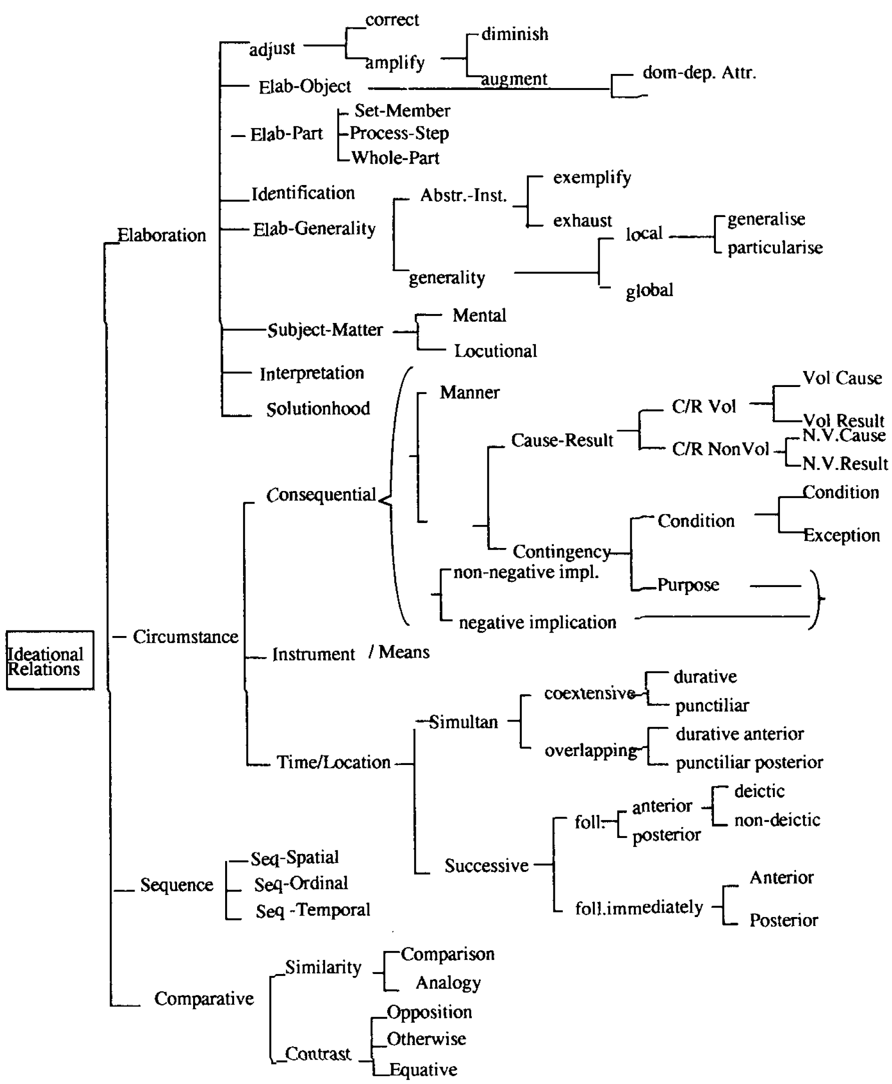

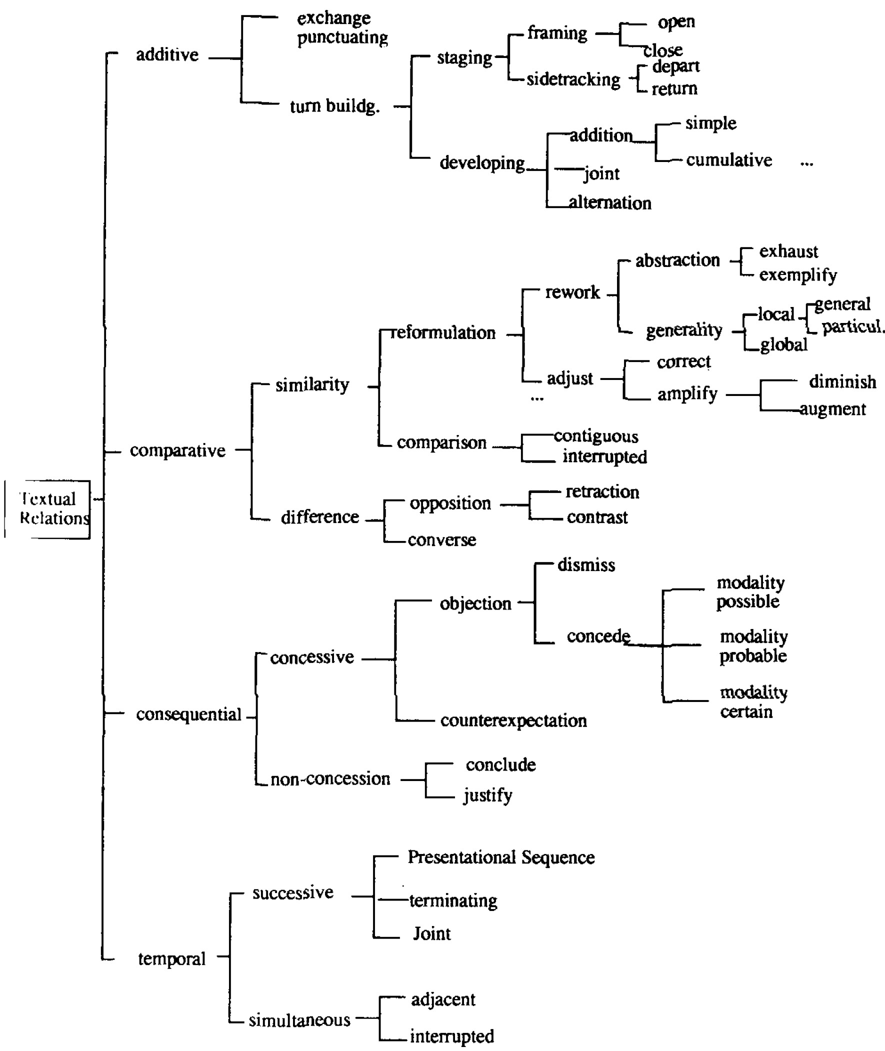
Textual Relations
additive
exchange
punctuating
turn buildg.
staging
framing
open
close
sidetracking
depart
return
developing
addition
simple
cumulative
...
joint
alternation
comparative
similarity
reformulation
rework
abstraction
exhaust
exemplify
generality
local
general
particul.
global
adjust
correct
amplify
diminish
augment
...
comparison
contiguous
interrupted
difference
opposition
retraction
contrast
converse
consequential
concessive
objection
dismiss
concede
modality possible
modality probable
modality certain
counterexpectation
non-concession
conclude
justify
temporal
successive
Presentational Sequence
terminating
Joint
simultaneous
adjacent
interrupted

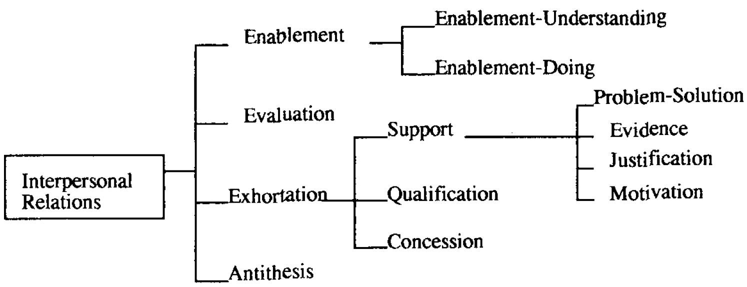

Figure-2- Text type Hierarchy

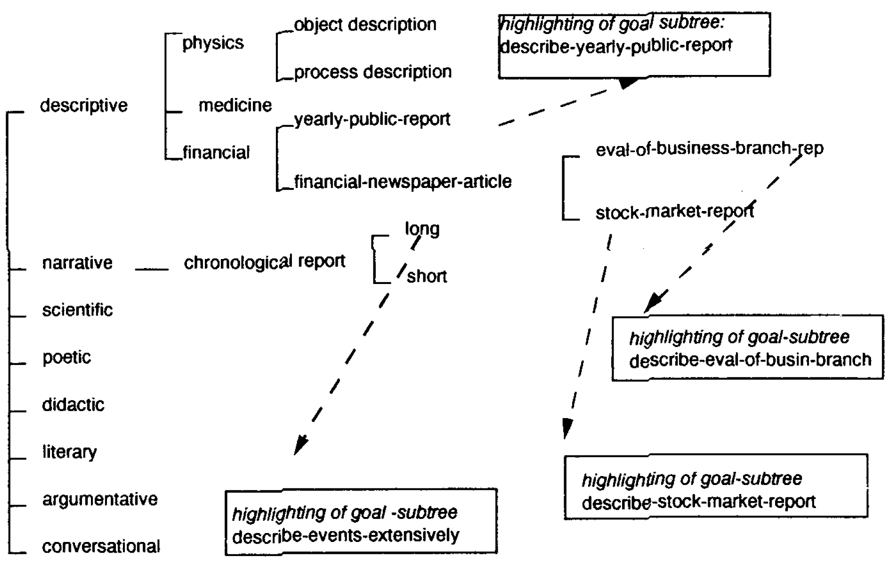

4.1.2. A text type network

This resource captures some aspects of the dimensions of variation in text generation based upon text type. The idea behind it is that different text types favour (highlight) different communicative goals and probably different discourse relations, although the motivating linguistic study is still to be done, and that once the user has specified a particular text type, the text planner will be constrained to use only certain relations and theme development patterns. Given its generality, the text type network to be used by the text planner is based on the hierarchy of text types proposed by de Beaugrande (1980), with extensions as needed to handle the text types we are addressing. The hierarchy, partially shown in Figure 2, is represented as a property-inheritance network in the knowledge representation system Loom.

4.1.3. A communicative goal network

This resource is responsible for capturing what Grosz and Sidner (1986) have called the intentional structure of discourse: a representation of the effect each part of the text is intended to have on the hearer and how the complete text achieves the overall discourse purpose (e.g., describe entity, persuade hearer to perform and action, etc.). At present, this resource in only rudimentary in the text planner; it is a rudimentary taxonomization of communicative goals, starting at the topmost level with some very general goals, such as INFORM , DESCRIBE , REQUEST and ORDER, which are eventually refined into specific goals to describe (or relate) specific types of information for specific contexts, see Figure 3. When complete, it will contain a definition of each communicative goal the system can recognize. Our taxonomy, which is implemented as a property-inheritance network, resembles the one being derived from Speech Acts by Allen (1991) and his colleagues.

Figure -3-. Hierarchy of Communicative Goals

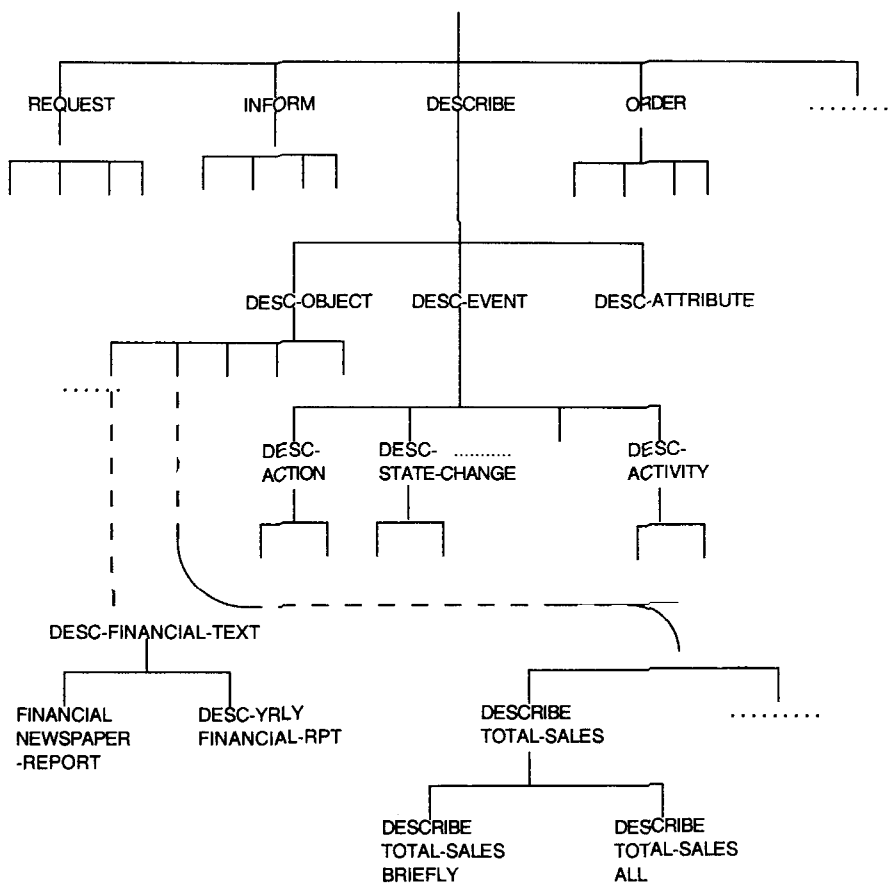

4.1.4. A generic structure potential library

Many genres contain some stereotypical texts —a business report, a financial newspaper article, an encyclopedia entry for an object, etc.— whose structure is not so much logically derived but passed along by custom. Such stereotypical structure has been studied by Hasan (1977; 1984) and represented as so-called Generic Structure Potentials (GSPs) (see also Martin 1991; Ventola 1987); in the AI and CL worlds, the notions of *script* (Schank and Abelson 1977) and *schema* (McKeown 1982), respectively, have the same function.

Our new text planner contains the beginnings of a GSP library, a library of stereotypical text structures. Each GSP is an ordered list of communicative goals, and each GSP serves to achieve some communicative goal. For example, a rather general GSP to achieve the goal "write-letter-to-friend" may contain, in sequence, the goals 1. "write our address"; 2. "write date"; 3. "write salutation"; 4. "write-thanks-for previous communication"; 5. "write news"; 6. "write closing"; 7. "sign name". Any of these goals may in turn be achieved by a more specific GSP, or if not then by a direct speech act such as *Inform* or *State*.

As a pilot experiment we chose a financial report as the target text to generate with the new text planner architecture; that is why the current GSP library is purely illustrative of the range of communicative goals available when generating that particular text. It is interesting to note that the specification of subgoals within a given overall discourse purpose reflects the text's generic structure as it is realised in stages.

4.1.5. Theme development

As we said before, the only textual resources that text planning systems generally used in the past were discourse structure relations and the pre-theoretical notion of focus. While the former are responsible for driving the text forward, the latter has been mainly used for tracking participants in discourse and for the resolution of anaphora. To date, no account has been given of other textual resources such as theme development, information distribution, or lexical cohesion.

However, theme fulfills a clear discourse function, and as such, it should have its place in a text planner design. The claim is that the selection of themes and the lexical material that is chosen to encode them in successive clauses in a text is not random, but guides what Fries (1983) has called the "method of development" of the text. Thus, themes are selected to highlight the points of rhetorical development, e.g., the points of elaboration, succession, and contrast. Theme reflects the structure of the passage and is, therefore, sensitive to a text's generic structure where it is realised in stages. Furthermore, thematic progression patterns seem to correlate with text type: in argumentative or expository prose, for example, the point of departure of each sentence should follow logically from what has gone before, while narratives favour the repetition of the same theme (normally the same participant) throughout successive sentences.

Though the study of theme has been traditionally restricted to the level of the sentence, its role is particularly noticeable at the clause-complex and discourse levels. This should be taken into consideration by a text generation system. Given a text to be generated, the system must establish how theme development may proceed and how themes are to be marked in each clause. The following three concerns arise:

- the type of theme to select; following Halliday (1985), there can be three different and simultaneous themes in each clause: the ideational, the interpersonal, and the textual.
- the theme progression pattern involved: the new theme can be the same as the theme of the previous clause; it may be part of the rheme of the previous clause; or it may be an element of what is called the "hypertheme" or general discourse segment topic (see Danes 1974).
- the linguistic degree of markedness of the theme: realization depends on the type of clause.

Figure 4 presents a portion of the theme network showing the selection of themes at the level of the clause.

4.1.6. Lexical cohesion

The idea of cohesion as a unity-creating "device" is well-known in linguistics (Halliday and Hasan 1976). If the lexical items in a text can be related to preceding or to following items through cohesive relationships, the text is seen more closely knit together. This is what is meant by texture. How does lexical cohesion function in a text? Lexical cohesion gives a text a "certain consistency of topic and predictability of development" (Halliday and Hasan 1976, 288). It guarantees that our discourse does not aimlessly wander from one discourse topic to another. As the text unfolds, what has preceded provides a context for later lexical items. In this sense, the study of lexical relations is not only interesting because it determines how well-constructed texts are, but also because the patterns of cohesion reveal something about the semiotic organization of texts (Ventola 1987). Lexical strings provide information about the relationship between the generic structure elements: a different type of lexical items takes over after each text stage has been realised; and these lexical items and the cohesive relations among themselves also reflect the rethorical development of the text.

This resource has not been implemented in the text planner yet. One of the main priorities when extending it will be to make this resource operative in the form of a network.

4.2. *The text planning process*

Having described the linguistic knowledge used in the text planner —knowledge encoded as static resources in a declarative format— we now turn to the other side of the coin, the process of planning, or procedural, dynamic resource. The purpose of the planning process is to

Figure 4. A Portion of the Theme Network

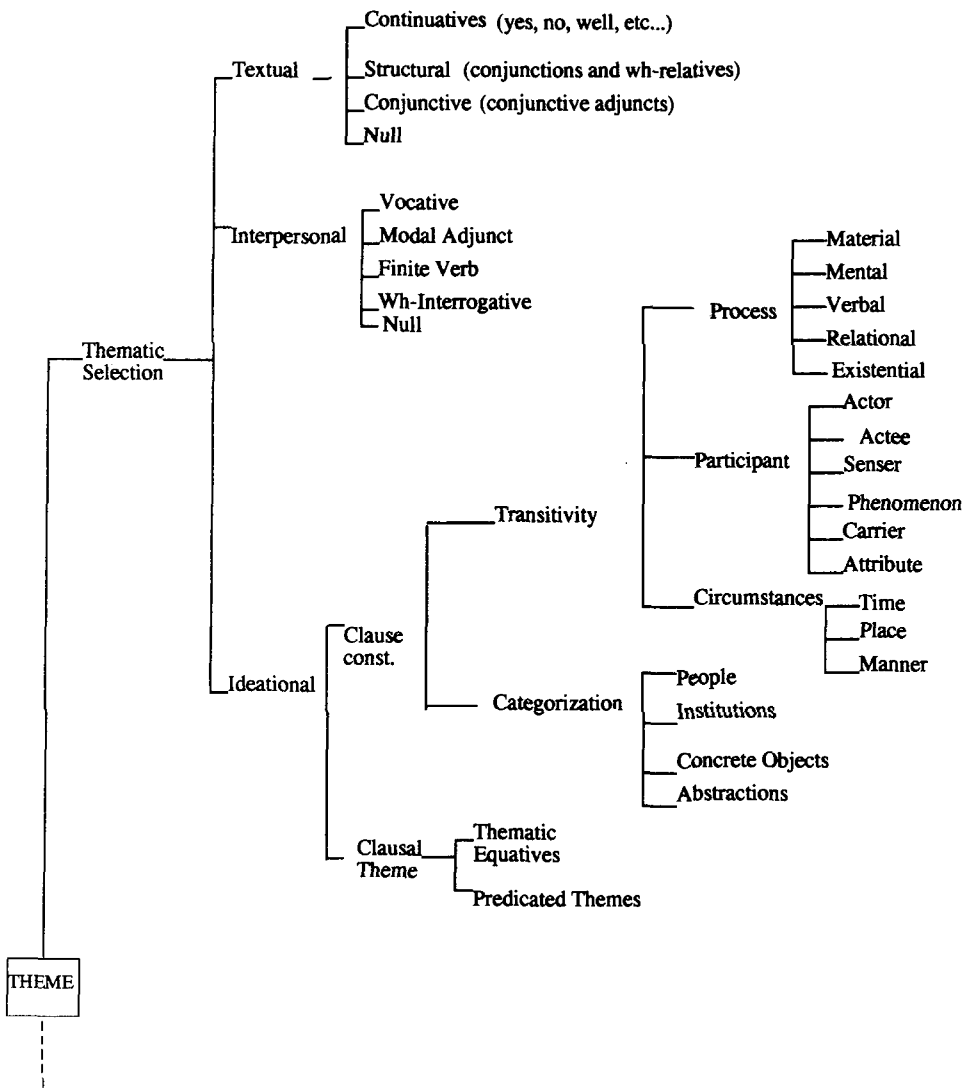

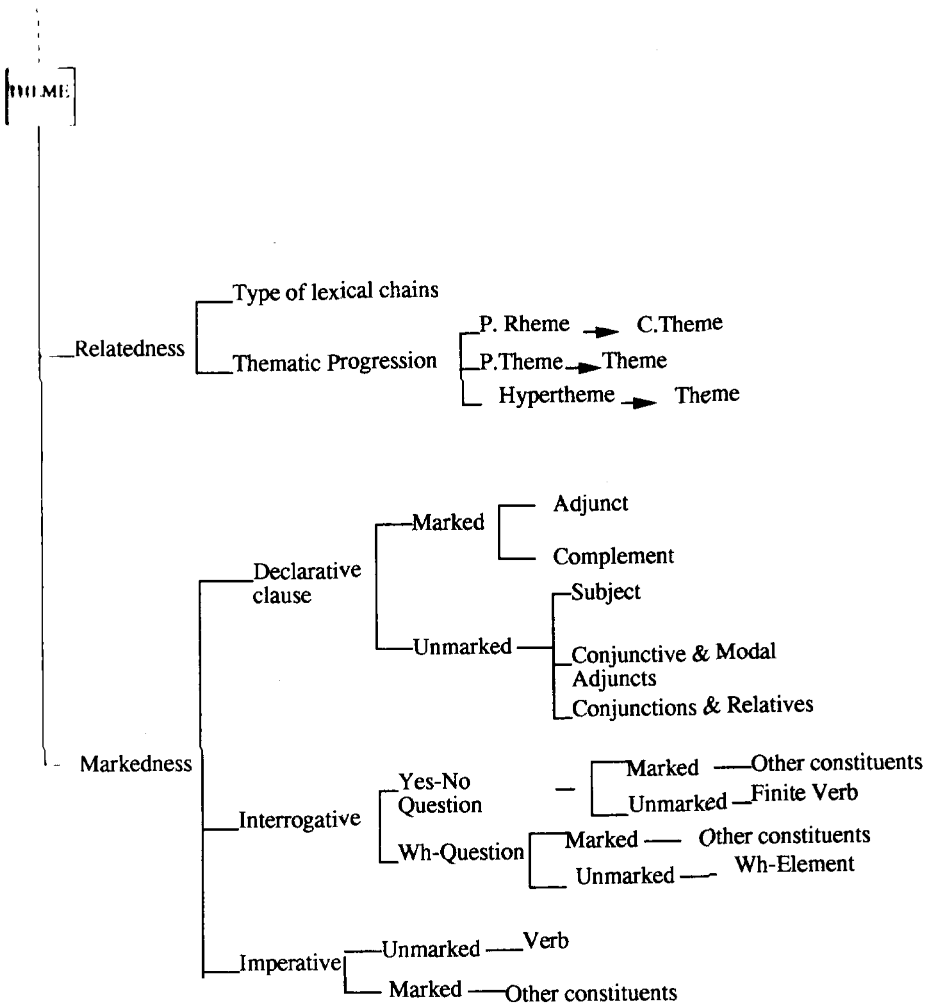
Relatedness
Type of lexical chains
Thematic Progression
P. Rheme → C.Theme
P.Theme → Theme
Hypertheme → Theme
Markedness
Declarative clause
Marked
Adjunct
Complement
Unmarked
Subject
Conjunctive & Modal Adjuncts
Conjunctions & Relatives
Interrogative
Yes-No Question
Marked — Other constituents
Unmarked — Finite Verb
Wh-Question
Marked — Other constituents
Unmarked — Wh-Element
Imperative
Unmarked — Verb
Marked — Other constituents

(1) find the next appropriate linguistic knowledge to apply;
(2) apply it in the current context;
(3) make appropriate changes in the context, either to the discourse structure tree being constructed, or to the general bookkeeping information required to run the planner;
(4) return to step 1, unless the planning process is complete; in which case continue with step 5;
(5) decompose the discourse structure tree into sentence-sized chunks and activate the Penman sentence generator to create an English sentence from each chunk.

The planner is initially activated with a communicative goal, such as "describe-a-financial-yearly-public-report" and one or more entities that represent (in a Knowledge Representation System called *Loom*) the most central piece of ideational content to be conveyed.

Typically, planning then proceeds as follows: after ratifying that the goal given is a well-defined goal, the planner searches in the GSP library for a GSP that matches the goal. If it finds one, it activates as subgoals each of the steps of the GSP, in order. It also creates a node in the fledgling Discourse Structure Tree with the name of the Goal and its GSP.

Having done this, the planner then tackles the first of the goals provided by the GSP. If this goal in turn provides another GSP, the subportions of the subordinate GSP are activated as goals and the subGSP is added to the Discourse Structure Tree. If, on the other hand, no GSP is found, then the planner activates any Discourse Relations in the Discourse Relations network that may apply. Each such Discourse Relation is treated in turn: the planner tests whether the relation can be applied to the current topic (the piece of ideational knowledge under consideration); if so, the relation provides an additional bit of knowledge from the Knowledge Base as satellite. The successful relation and the two pieces of knowledge are added to the bottom of the Discourse Structure Tree, under the goal that gave rise to them, and planning continues with the next Discourse Structure Relation or the next GSP goal, depending on what is next on the planner's agenda. Planning continues, with the

Discourse Structure Tree growing ever more bushy, until no more goals or relations are found to apply.

Of course, this is a rather simplified version of the planning process. In reality, several additional actions are performed before any Discourse Goal or Discourse Relation is successfully added to the Tree: themes are specified, ideational content is inspected, relations are blocked from view or are preferred, etc. In addition, considerable bookkeeping is required to keep all the planner's internal variables and data structures (the agenda, etc.) in order.

However, the basic idea is straightforward. The planning process is activated by a specific goal or a specific piece of information, and by accessing appropriate information from its different linguistic resources, the planner gradually constructs a Discourse Structure Tree that captures the structure of the text and organizes all the information that has been found. When done, the planner activates the sentence generator on each portion of the tree to construct a coherent English text.

5. Conclusion

My aim in this paper has been to introduce to the linguistic community some preliminary findings in the area of text planning which I hope will serve to lay the foundations of a computational theory of discourse structure. I have also argued why a text planning system must separate linguistic knowledge from the control mechanism that activates and interconnects those resources, and have presented the way in which those two aspects complement each other. The current text planner we are developing aims at bridging the gap in previous text planning systems where much of the linguistic information is not explicitly represented declaratively and still resides as part of the control mechanism. I have presented some of those resources which have been successfully identified and represented to date, and the planning process that activates them. This task, far from complete, is, in my view, a stimulating research arena both for linguists and computer scientists. It is my hope that this presentation encourages others to follow this avenue of research so that both scientific communities can benefit from the interaction.

WORKS CITED

Allen, J. F. 1991. "Discourse Structure in the TRAINS Project". *Proceedings of the DARPA Speech and Natural Language Workshop.* San Mateo, CA: Morgan Kaufmann.

Appelt, D. E. 1985. *Planning English Sentences.* Cambridge: Cambridge U.P.

Beaugrande, Robert de. 1980. *Text, Discourse, and Process: Toward a Multidisciplinary Science of Texts.* London: Longman.

Carberry, S. M. 1983. "Tracking User Goals in an Information-Seeking Environment". *Proceedings of the National Conference on Artificial Intelligence* AAAI-83, Washington D.C.

Cawsey, A. 1990. "Generating Communicative Discourse". R. Dale, C. Mellish and M. Zock. *Current Research in Natural Language Generation.* Eds. R. Dale, C. Mellish and M. Zock. New York: Academic.

Clippinger, J. H. 1974. *A Discourse Speaking Program as a Preliminary Theory of Discourse Behaviour and a Limited Theory of Psychoanalytic Discourse.* PhD Diss. University of Pennsylvania.

Cohen, P. R. 1978. *On Knowing What to Say: Planning Speech Acts.* PhD Diss. University of Toronto.

Danes, F. 1974. "Functional Sentence Perspective and the Organization of the Text". *Papers on Functional Sentence Perspective.* Ed. F. Danes. The Hague: Mouton.

Emele, M. et al. 1990. "Organizing Linguistic Knowledge for Multilingual Generation". *Proceedings of the 13th International Conference on Computational Linguistics.* Helsinki.

——. 1991. "Interactions between Linguistic Constraints: Procedural vs Declarative Approaches". *Journal of Machine Translation, special edition.*

Fries, P. H. 1983 "On the Status of Theme in English: Arguments from Discourse". *Forum Linguisticum* 6: 1-38.

Garnham, A., J. V. Oakhill and P.N. Johnson-Laird. 1982. "Referential Continuity and the Coherence of Discourse". *Cognition* 11: 29-46.

Givón, T. 1983. *Topic Continuity in Discourse: Quantitative Cross-Language Studies.* Typological Studies in Language, vol. 3, Amsterdam: Benjamins.

Grimes, J. E. 1975. *The Thread of Discourse.* The Hague: Mouton.

Grosz, B. J., and C. L. Sidner. 1986. "Attentions, Intentions and the Structure of Discourse". *Computational Linguistics* 12: 175-204.

Halliday, M. A. K., and R. Hasan. 1976. *Cohesion in English*. London: Longman.

Halliday, M. A. K. 1985. *Introduction to Functional Grammar*. London: Arnold.

Hasan, R. 1977. "Text in the Systemic Functional Model". *Current Trends in Textlinguistics*. Ed. W. Dressler. Berlin: de Gruyter.

Hasan, R. 1984. "The Nursery Tale as a Genre". *Nottingham Linguistic Circular* 13: 71-102.

Hobbs, J. R. 1978. "Why is Discourse Coherent?" *SRI Technical Note 176*.

. 1979. "Coherence and Coreference". *Cognitive Science* 3: 67-90.

. 1982. "Coherence in Discourse". *Strategies for Natural Language Processing*. Eds. W. G. Lehnert and M. H. Ringle. Hillsdale, NJ: Erlbaum.

Hovy, E. H. 1988a. "Planning Coherent Multisentential Text". *Proceedings of the 26th ACL Conference*, Buffalo.

. 1988b. "Approaches to the Planning of Coherent Text". *Natural Language in Artificial Intelligence and Computational Linguistics*. Eds. C. L. Paris, W. R. Swartout and W. C. Mann. Boston, MA: Kluwer Academic Publishers.

—. 1988c. *Generating Natural Language under Pragmatic Constraints*. Hillsdale, NJ: Erlbaum.

—. 1990. "Parsimonious and Profligate Approaches to the Question of Discourse Structure Relations". *Proceedings of the 5th International Workshop on Language Generation*. Pittsburgh, PA.

Jameson, A. 1987. "How to Appear to be Conforming to the 'Maxims' even if You Prefer to Violate Them". *Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology and Linguistics*. Ed. G. Kempen. Boston, MA: Kluwer Academic Publishers.

Kintsch, W., and T. A. van Dijk. 1978. "Toward a Model of Text Comprehension and Production". *Psychological Review* 85: 363-94.

Linde, C. 1979. "Focus of Attention and the Choice of Pronouns in Discourse". *Syntax and Semantics*. Ed. T. Givón. Vol. 12. New York: Academic.

Maier, E., and E. H. Hovy. 1991. "A Metafunctionally Motivated Taxonomy for Discourse Structure Relations". *Proceedings of the 3rd European Workshop on Language Generation*. Judenstein.

McKeown, K. R. 1982. *Generating Natural Language in Response to Questions about Database Queries*. PhD Diss. University of Pennsylvania.

Mann, W. C. 1983. "A Linguistic Overview of the Nigel Text Generation Grammar". *Proceedings of the Xth International LACUS Forum*. Quebec.

Mann, W. C., and S. A. Thompson. 1983. "Relational Propositions in Discourse". USC/ ISI Research Report RR-83-115.

——. 1987. "Rhetorical Structure Theory: Description and Construction of Text Structures". *Natural Language Generation: New Results in Artificial Intelligence, Psychology, and Linguistics.* Ed. G. Kempen. Dordrecht: Nijhoff. Also USC/ ISI Research Report RR-86-174.

Martin, J. R. 1991. *English Text: System and Structure.* Amsterdam: Benjamins.

Matthiessen, C. 1987 "Notes on the Organization of the Environment of a Text Generation Grammar". *Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology and Linguistics.* Ed. G. Kempen. Boston, MA: Kluwer Academic Publishers.

——, and J. A.Bateman. 1991. *Text Generation and Systemic-Functional Linguistics: Experiences from English and Japanese.* London: Pinter.

Maybury, M. T. 1990. *Planning Multisentential English Texts Using Communicative Acts.* PhD Diss. University of Cambridge. Available as In-House Report TR-90-411, RADC, Rome, NY.

Meyer, B. J. F. 1975. *The Organization of Prose and its Effects on Memory.* Amsterdam: North Holland.

Moore, J. 1989. *A Reactive Approach to Explanation.* PhD. Diss. UCLA.

Paris, C. 1987. *The Use of Explicit User Models in Text Generation: Tailoring to a User's Level of Expertise.* PhD Diss. Columbia University.

——. 1988. "Tailoring Object Descriptions to the User's Level of Expertise" *Computational Linguistics* 14: 64-78. Special Issue on User Modeling.

The Penman Documentation. 1989. Unpublished documentation for the Penman Language generation system, USC / ISI.

Polanyi, L. 1988. "A Formal Model of the Structure of Discourse". *Journal of Pragmatics* 12: 601-38.

Redeker, G. 1990. "Ideational and Pragmatic Markers of Discourse Structure". *Journal of Pragmatics* 14: 367-81.

Reichman, R. 1981. "Conversational Coherence". *Cognitive Science* 2:283-327.

Reichman-Adar, R. 1984. *Plain Speaking: A Theory and Grammar of Spontaneous Discourse.* PhD Diss. Harvard Unversity. BBN Report No. 4681.

Sanders, T., and W. Spooren. 1991. "Organizing the Set of Coherence Relations" Presented at the *Workshop on Discourse Coherence.* University of Edinburgh.

Schank, R. C. and Abelson, R. P. 1977. *Scripts, Plans, Goals and Understanding.* Hillsdale, NJ: Erlbaum.

Sidner, C. L. 1983. "Focusing and Discourse". *Discourse Processes* 6: 107-30.

Swartout, W. R. 1983. "XPLAIN: A System for Creating and Explaining Expert

Consulting Systems". *Artificial Intelligence* 21: 285-325. Also USC/ISI Research Report Rs-83-4.

Ventola, E. 1987. *The Structure of Social Interaction: A Systemic Approach to the Semiotics of Service Encounter Interaction.* London: Pinter.

Woolf, B. P. and McDonald, D. D. 1984. "Context Dependent Transitions in Tutoring Discourse". *Proceedings of the 4th AAAI Conference.* Austin.

INDICE