Natural Language and Terminology

Strengthening Argumentation in Medical Explanations by Text Plan Revision*

Fiorella de Rosis[1], Floriana Grasso[2] and Dianne C. Berry[3]

[1] Dipartimento di Informatica, Università di Bari, Italy
e-mail: derosis@gauss.uniba.it
[2] Department of Computing & Electrical Engineering, Heriot-Watt University, Edinburgh, UK, e-mail: floriana@cee.hw.ac.uk
[3] Department of Psychology, University of Reading, UK
e-mail: d.c.berry@reading.ac.uk

Abstract. Explanations are an important by-product of medical decision-support activities, as they have proved to favour compliance and correct treatment performance. To achieve this purpose, these texts should have a strong argumentation content and should adapt to emotional, as well as to rational attitudes of the Addressee. This paper describes how *Rhetorical Sentence Planning* can contribute to this aim: the rule-based plan discourse revision is introduced between Text Planning and Linguistic Realization, and exploits knowledge about the user personality and emotions and about the potential impact of domain items on user compliance and memory recall. The proposed approach originates from analytical and empirical evaluation studies of computer generated explanation texts in the domain of drug prescription.

1 Introduction

Explanations are an important by-product of any medical decision support activity, as they substantiate the system suggestions by favouring, at the same time, compliance and correct treatment performance. To achieve their goals, these messages have to be formulated in a clear and convincing way, and their arguments have to be adapted to the addressee. Whenever possible, they have to be generated by exploiting information from already established sources, rather than requiring ad hoc data collection.

In the last few years, several text generation systems have been designed in order to produce effective explanations. In these systems, the generation process is made up of a phase of *Text Planning*, in which information content and order and rhetorical structure are established, followed by a phase of *Linguistic Realisation*, in which the plan is translated into an understandable, coherent message [7]. More recently, the limitations of this approach have become apparent, and several authors are now investigating how to bridge the gap between the

* This work was partially supported by a British-Italian Collaboration in Research and Higher Education Project, which involved the Universities of Reading and of Bari, in 1996.

two phases. Methods proposed introduce, between Text Planning and Linguistic Realisation, a phase which is usually denoted as *Sentence Planning* [8, 10]. In this intermediate phase, *local* refinement techniques are applied to increase the quality of the generated text. For example: aggregating clauses or elements internal to a clause, defining how to realise rhetorical relations, or the salience nucleus/satellite and so on [11, 2, 13].

It is our opinion that, although there is a real need for style refinement in complex text generation (especially when existing information sources are employed), this new phase of *Plan Revision* should also be aimed at strengthening argumentation in the text: this requires developing methods which consider the emotional state of the addressees, as well as their knowledge. This claim is indirectly supported by Dalianis and Hovy: "text length (i.e. redundancy of words) is not the best measure of readability of aggregated texts. Instead, a better measure is internal (structural) coherence" [2]. Sentence optimisation then originates not only from linguistic needs, but also (in some cases, mainly) from rhetorical ones.

On the other side, the importance of argumentation in medical explanations has been proved by several experimental psychologists, who have argued that "enlistment of the patient by the physician is promoted by the use of inclusive language, that is, language that explicitly recognises patients' competence and thus treats them with respect. We believe that such language increases the physician's (and possibly also the prospective system's) chances of having the patient hear the diagnosis and treatment recommendations" [5].

In her paper about the generation of *empathetic responses* in a dialogue about medical treatment, Haimovitz [6] argues that, to produce texts suited to the direct and the indirect users' needs, sentences with a given information content should be generated so as to "stress favourable information while downplaying or offsetting unfavourable information". This effect can be obtained by exploiting knowledge about the indirect user's *concerns* and *worries*, in addition to domain knowledge about the possible *impact* of information items. A *very sensitive test* or a *non invasive treatment* are examples of favourable information, whereas *a very long hospital stay* or *a painful procedure* are examples of an unfavourable one. The *cost* of an utterance is calculated as a function of the addressee's characteristics (for example: anxiety about pain); the *accent* of a sentence is then modified by joining sentences with appropriate conjunctions, or by introducing *detensifier* or *intensifier* adverbs.

The work reported in this paper is the continuation of a European Research Project OPADE which was aimed at designing and prototyping a decision support system in the domain of drug prescription. A text generator was designed as part of that Project, from a corpus of explanations provided by doctors. Texts generated were submitted to informal and formal evaluations, which confirmed their interest but also revealed a number of limitations in the method developed. In this paper, after briefly describing the structure of our generator (Sect. 2), we will examine these limitations (Sect. 3). We will then prove that they cannot be solved by a purely stylistic revision of the text, but require a more complex *Plan Revision and Refinement*, which exploits knowledge about the domain, as

well as about the addressee (Sect. 4). We will then outline (in Sect. 5) a method which solves part of these problems. Some consideration about the limitations of this approach will conclude the paper (Sect. 6).

2 The Context

The explanation component of OPADE generates a text which illustrates the prescription by retrieving information from a set of heterogeneous databases: a medical record, a drug database and a prescription record. The generation process combines hierarchical Text Planning with a Linguistic Realisation based on Augmented Transition Networks [3]. The message produced is adapted to the Speaker (the doctor who made the prescription) and to the addressee (the patient, a nurse or a doctor's colleague): two user models represent the main attitudes of the two users. Adaptation of the message is made in the planning phase (by introducing in plan operators conditions about the context), as well as in the Linguistic Realisation (by attaching similar conditions to ATN's arcs). The discourse plan is represented in OPADE as a tree, whose leaves correspond to primitive communicative actions; a communicative goal and a rhetorical relation are associated with its intermediate and root nodes. An example of the top part of a plan tree is shown in Fig. 1. This plan includes the following main

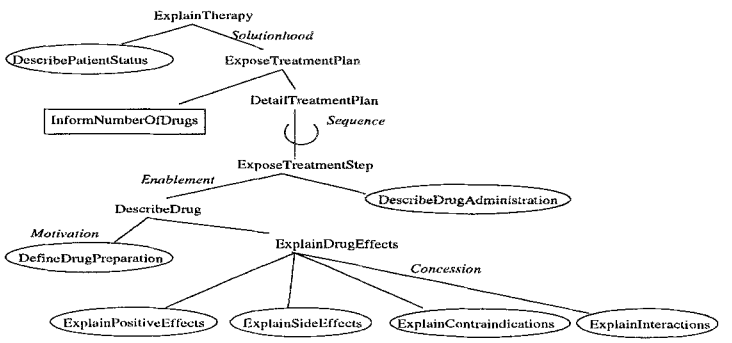

Fig. 1. An example of the top part of a discourse plan

components (which are encircled with an oval):
• A *Persuade about the need of treatment*, realised by means of a description of the health problems that the prescription intends to solve.
• A *Request to perform the treatment*, with information about the number of drugs in the prescription.
• A detailed description of the treatment plan. Drugs are considered in sequence, by introducing, for each of them: (i) a *Request to administer the drug*, with a description of its characteristics; (ii) a *Persuade that the drug is useful and not harmful*, with an illustration of its positive effects (efficacy) and of the negative

ones (side effects, contraindications, interactions with other drugs), and (iii) an *Enable to administer the drug correctly*, with a detailed description of administration modalities.
An example of text generated from this plan in shown in Fig. 2. This text is subdivided (for explicative purposes) into *discourse segments* which are the expansion of the encircled goals in Fig. 1. A sample of texts generated was evaluated by

Comments to the drug prescription for Mr Fictif

DS1 You have been diagnosed as suffering from a mild form of what we call 'angina', that is a spasm of chest resulting form over-exertion when hearth is diseased. In addition, you have elevated cholesterol.
DS2 To solve these problems, there are two drugs I would like you to take.
DS3 The first one is Aspirin, which is a analgesic, that is it relieves the pain.
DS4 This drug has been shown to reduce the risk of hearth attack and prevents abnormal blood clotting within five minutes.
DS5 The only problem is that this drug can be associated with some side effects. The first one is bleedings, that is a slight hemorrhage; it can be a serious side effect but occurs infrequently. In the unlikely event that this occurs, I suggest you to stop treatment. The second one is allergic reaction; it can be a serious side effect but doesn't happen commonly; it occurs especially when the patient is suffering from what we call 'atopic disease', that is a hypersensitivity to substances having a basis of hereditary predisposition. In case you notice that this problem occurs, you should stop treatment.
DS6 I must also prevent you that this drug may interfere with many others. One interaction is with Vitamin K inhibitors; this could cause increased bleeding risk. Therefore, please follow carefully your prescription, do not take other drugs without telling your physician and do not stop treatment without your physician's advice.
DS7 You have to take this drug by mouth; take it with water, at regular intervals. If you miss a dose, do not modify rhythm or quantity of the following ones. Regarding administration of this drug, the treatment steps are the following. The first one is a attack treatment. The dosage is 1 gr in 2 intakes for 1 day. You have to take one tablet of 500 mg at 8 am and one tablet of 500 mg at 8 pm. The second one is a maintenance treatment. The dosage is 500 mg in 1 intake for 7 days. You have to take one tablet of 500 mg at 20.
DS8 The second drug I will give you is Glyceril Trinitrate. . .
. . . . . .
DSf These are the main treatment suggestions I wanted to give you. If you need any more information, please do not hesitate to ask me directly.

Fig. 2. An example of text generated from OPADE

representatives of OPADE's final users: in this informal evaluation, explanation texts were considered as very interesting for both patients and health professionals, especially considering that, in general, this type of information is only provided verbally (if at all). Subsequently, we performed two more formal evaluation studies. In a first *analytical evaluation*, we compared the artificial texts with the corpus of explanations that were employed to define the generation method. In a second, *empirical study*, we evaluated the efficacy, on the addressee, of texts built by varying information content and order. The two studies highlighted several limitations in the generated texts, and provided a number of cues on how better texts might be obtained.

3 Limitations of Texts Generated

3.1 Comparison with the Corpus of Explanations

Our corpus of explanations included 36 texts, produced by 6 doctors on 4 clinical cases (two cases of angina in a male and a female aged patient, a case of tuberculosis and a case of subacute thyroditis). Although the computer generated explanations reproduce the overall characteristics of these texts, several elements contribute to increase the *argumentative strength* of natural messages, and are not reproduced in the artificial ones.

Enlistening Techniques in the 'Persuade about the need of treatment'. This goal is obtained by linking the two discourse segments *Describe Patient Status* and *Expose Treatment Plan* by a rhetorical relation (RR) of Solutionhood. This RR may become a Concession when the patient is not seriously ill, conveying the idea that, although the disease mildness might let one presume that treatment is not needed, this is not the case:

Ex1: This condition is mild, *but* we have to treat it....

On the other side, the Description of the patient's health status contains a RR of Elaboration Object-Attribute, whose nucleus is the name of the diagnosis and whose satellites are its severity and certainty. The value of these satellites influences the nucleus description. For example:

Ex2: *Unfortunately* you have an infection..., (when the severity is *high*).

Persuasion Elements in the 'Request to perform the treatment'. The goal *Request to perform the treatment* is achieved, in our generator, by a simple sentence which indicates the number of drugs in the prescription. In the corpus of texts examined, elements of persuasion are introduced, in this segment, by means of several techniques:

(1.) To predispose favourably the addressee towards the Request: by anticipating the positive aspects of treatment, when it has a high expected effectiveness or when it is, at least, expected to have a positive effect on symptoms:

Ex3: *So, the good news* is that we do have tablets that are very effective against TB...

Ex4: With the treatment we'll give you, *we can bring all these things back to normal*...

Ex5: We can give you something *to make you feel more comfortable*...

(2.) To prevent a non collaborative attitude: by anticipating aspects on which a high compliance is needed, especially when treatment is long and complex:

Ex6: We have to undertake *quite a long course of treatment*...

(3.) To promote remembering of significant items: by synthesising information items which are common to several drugs, *before* detailing the treatment:

Ex7: ... two drugs which are very similar: *they both aim to open the blood vessels up.*

Ex8: I'm going to ask you to take *two tablets*...

Ex9: ... *some tablets* to take *twice a day*...

Ex10: ... and you will have to take them *for some months.*

or by means of a *final synthesis* of this detailed description:

Ex11: So, in summary, we'll give you *tablets* to *help the spasm of the blood vessels*...

(4.) Persuasion elements in the 'Request to administer the drug'. This purpose is achieved by associating treatment to procedures or drugs with which patients are familiar and that they know to be *harmless*:

Ex12: The drug that I'm going to ask you to take is Aspirin, *just ordinary Aspirin*...

or by emphasising positive aspects of treatment:

Ex13: ... one *small* dose of...

(5.) Strengthening or attenuation techniques in the 'Persuade that the drug is not harmful'. In this case, positive aspects are overemphasised by means of particular linguistic markers and negative aspects are offset by using *aggregation* techniques to avoid repetitions :

Ex14: *The only possible side effect* this may have is to...

Ex15: *Again*, it's very unlikely to happen; *Again*, I would like you to reassure that...

Ex16: *Occasionally*, they can disturb you in your sleep.

3.2 Cues from Evaluation Studies

At the University of Reading, several studies were made to investigate whether and how the order of presentation of information and the level of detail influence the effectiveness of explanation, measured in terms of how good the explanation is, perceived likelihood of taking the medication and memory of information. To this aim, several texts were prepared, by combining differently the following components of an OPADE explanation: description of the patient health status, prescription, description of drug, side effects and contraindications. Experiments were repeated by varying number and seriousness of side effects. The results of these studies are described elsewhere [1]: here, we refer only those of them which are relevant to this paper.
The first finding is that the optimal order of presentation of information cannot be established in a rigid way, but depends on the overall explanation content. In particular, participants in the study overall recalled more information about drug administration when it was presented before information about side effects. However, when the explanation described only a few side effects, participants remembered more information about side effects if the information was presented before, rather than after, the information about drug administration. There was no effect of order when the explanation described a large number of side effects. Therefore the optimal order of information within the explanation can only be determined once the number of side effects has been established.
The second finding concerns the perceived likelihood of taking the medication and the evaluation of how good the explanation is, as a function of side effects' seriousness: if the text includes *negative information*, the need for additional explanatory information increases. Thus the decision about the level of detail of the explanation is a function of the overall information content: details have to increase according to the *negativity* of information elements in the text, such as the severity of the disease or of the side effects of drugs.

4 Some Hints on How to Solve the Mentioned Problems

We identify two limitations in current text generation methods, that reduce argumentation strength in the texts produced: (i) Text Planning is a one shot process which does not take into account the impact that the text 'as a whole' will have on the addressee; (ii) the subsequent phases disregard the rhetorical nature of the text. Let us justify more deeply our claim.

4.1 Text Planning Shortcomings

In hierarchical Text Planning, information content and presentation order are established as a function of the main communicative goal, of hypotheses about the mental states of the Speaker and the addressee and of domain knowledge.

In their decision on whether to apply a specific plan operator, top-down plan expansion methods consider only the values of information items which relate to the discourse segment which is being planned. The overall content of the text is not known until planning has been completed. For example: when the subtree concerning one of the side effects of a specific drug is being expanded, the planner doesn't yet *know* how many side effects the drug will have and whether these side effects will be serious, frequent and so on. The discourse structure and its content will be established entirely only when the plan will have been completed: problems mentioned in Sect. 3 will only then become manifest.

4.2 Sentence Planning Shortcomings

As we mentioned in the Introduction, the need for an intermediate phase between Text Planning and Linguistic Realisation is acknowledged by most researchers in the domain, and the tasks this phase is supposed to undertake have been specified by several authors. Among them, a fine grained local discourse structuring, which includes decision about inclusion/exclusion of discourse markers, a sentence content delimitation, an internal sentence organisation, with aggregation, reference and lexical choice [8]. Concepts involved in this phase are essentially linguistic-stylistic ones: focus, pronominalisation, preposition phrases, etc.
We think that, in most of these tasks, rhetorical concepts should be considered as well, and that other rhetoric-driven tree-restructuring tasks need to be performed. For example:

1. adding empathy to sentences: by stressing positive concepts (Ex12,13,14), offsetting negative ones (Ex16), marking a non-aggregation choice (*again* in Ex15) or showing an empathetic or enlistening attitude (Ex2).

2. treatment of repetitions: *aggregation* is defined as the process of removing redundant information in a text without loosing any information [2]. The aim of this process is to obtain more fluent and easy to read texts. So from the three sentences: *Tom loves Mary, Dick loves Mary, Harry loves Mary*, aggregation will produce the sentence: *Tom, Dick and Harry love Mary*, which is much more fluent, and undoubtedly natural. But, let us suppose that this was an argumentative text, aimed at convincing the addressee of the kindness of Mary. The text will be much more effective if repetition is employed (and possibly emphasised with an emphatic punctuation): *Tom loves Mary! Dick loves Mary! Harry loves Mary!* which means: *Everybody loves Mary, who couldn't?*
The decision on whether to aggregate repetitions or to emphasise them should therefore be made by considering the global aim of the sentence, not only its readability. Ex15 reveals that doctors are guided by this concept in their decisions. In addition, rhetorical aspects can help to decide, when aggregating, the order in which items should be presented to maximise the argumentative effect.

3. introducing opening or closing summaries: Examples from 3 to 11 showed how these summaries can be employed to predispose favourably the addressee or to promote memory of relevant facts.

4. changing information ordering: the result of studies by psychologists at

Reading showed that the information ordering established by the planner has sometimes to be revised, according to the overall discourse content; the same applies to the next item:

5. revising information detail: (see findings about side effects).

6. changing rhetorical relations: Ex1 is a typical case in which a rhetorical relation of Solutionhood can be less effective than a Concession when addressees have to be persuaded to undertake an action aimed at solving a problem that they perceive as minor.

Some of these tasks require a *Plan-Tree Refinement*, in which the main decisions of the plan (information order and content, RRs, communicative goals) are respected: we name this class of tasks *Rhetorical Sentence Planning*, to contrast them with current Sentence Planning techniques.
Another class of tasks introduces deeper changes in the plan structure and could be called *Plan-Tree Revision.*

5 Some Implementation Attempts

Rhetorical Sentence Planning and plan revision tasks require **upgrading knowledge** that is commonly employed for text generation. The Addressee mode, which usually includes general characteristics (eg. age, family status) and a large variety of *rational* attitudes (knowledge, beliefs, goals) is extended with domain-related, emotional aspects of personality, such as *hostile towards treatment* or *anxious*. This broadening of the user modelling sphere of interest corresponds to a trend in the domain of human-computer interaction, which claims that extrarational aspects of human beings should be considered in designing *believable* computer systems [4, 12]. Domain knowledge has to be upgraded as well, by attaching to the data items (object-attribute-value triples) one or more labels defining the semantic properties which may be relevant in the domain (eg. *favourable for compliance, relevant for correct treatment* and so on).
Both Rhetorical Sentence Planning and Plan Revision make a wide use of these knowledge sources, and can be represented as sets of tree-rewriting rules which apply to discourse subplans. For each rule:
The left-hand side is a logical combination of conditions on the communication goal and the RR associated with the root, the semantic properties of information associated to the leaves and the addressee's attitudes.
Examples of semantic properties of information are the following:
- *unfavourable* item: Severity (?disease, serious)
- *favourable* item: Efficacy (?drug, high)
- *relevant-for-compliance* item: Duration (?treatment, short)
- *relevant-for-correct-treatment* item: AdministrationForm (?drug, tablets).

Examples of addressee's attitudes:
- Compliance (?drug-treatment, low)
- Anxiety (?pain, high).

The right-hand side is a tree-restructuring algorithm which varies according to the task to be performed.

5.1 Some Examples of Rhetorical Sentence Planning Rules

Adding Empathy to Sentences.

R1 IF the subtree 'Persuade that the drug is not harmful' includes sub-subtrees whose leaves are 'favourable' items with equal values AND the patient is pain-anxious THEN emphasise these leaves.

The emphatization algorithm depends, in this case, on the Rhetorical Relation associated to the root node: it may consist of adding intensifier or detensifier adverbs or adjectives or in generating compound sentences in which the components are joined by a contrasting conjunction (as in [6]). Operations of this type may solve problems like those in Ex2 and Ex13, Sect. 3. In the first of them, a RR of Elaboration Object-Attribute is associated with the root. In this case, the two *Inform* are merged into a unique sentence, which is introduced by an appropriate expression: an intensifier (*the good news is that*) in case of favourable information, a detensifier (*unfortunately*) in case of unfavourable information. A similar transformation is applied to the second tree, by introducing the detensifier adjective *small*. These transformations are adapted to the addressee and exploit a semantic knowledge of the domain.

Rhetorical Aggregation.

R2 IF the subtree 'Persuade that the drug is not harmful' includes sub-subtrees whose leaves are unfavourable items with equal values AND the patient is pain-anxious THEN aggregate repetitions of unfavourable items in these subtrees.

This algorithm is applied to subtrees whose root-node is associated with an Ordinal Sequence or a Logical RR: repetitions in the text usually originate from these subtrees. To illustrate it, let us consider a segment of an OPADE's text:
However, I must inform you that this drug may cause some side effects. The first one is nausea; it is serious, it occurs infrequently, in a strong form, in sensitive patients. The second one is headache; it is serious, it occurs infrequently, in a strong form, in sensitive patients. The third one is insomnia; it is not serious, it occurs frequently, in a strong form, in sensitive patients....

The algorithm is made up of four main steps:

step 1: finding subtrees such that the RR associated to the root is an Ordinal Sequence or a Logical. In our example, this step extracts, from the plan tree in Fig. 1, the SubTree shown in Fig. 3. This subtree is made up of several *similar* (in their structure) sub-subtrees;

step 2: classifying leaves. The classification is made in two stages:

- identification of the *main nucleus* of each sub-subtree (according to the focusing structure) and labelling of these leaves as *objects*; labelling of remaining leaves as *characteristics*.

In our example, there are three nuclei (one in each element of the Sequence), corresponding to nausea, headache and insomnia; and four characteristics for each of them: severity, frequency, intensity and risk-category.

- classification of objects according to the values of their characteristics, and identification of characteristics whose values are the same in all objects (uniform

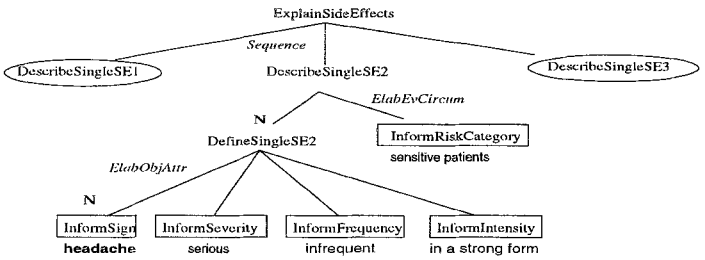

Fig. 3. A portion of the plan tree to explain side effects of a drug

characteristics) or only within a class (class-uniform characteristics).
In our case, frequency and severity are class-uniform characteristics; risk-category and action are uniform characteristics.

step 3: deciding which characteristics to aggregate.
In our example, *strong* intensity is an unfavourable item, *infrequent* frequency is a favourable one; the first one is therefore aggregated, the second one is duplicated.

step 4: restructuring the subtree:
4.1: prune uniform characteristics and add a new root and a new child leaf of this root for each of them;
4.2: rearrange the nuclei in the original Ordinal Sequence according to classes of *similar* objects; for each class which contains more than one object, create a new intermediate node and define an Ordinal Sequence among objects within it;
4.3: apply recursively 4.1 and 4.2 within each class.

Figure 4 shows how the subtree in Fig. 3 is transformed by this algorithm. If some empathy is added to sentences, the following text will be generated:
However, I must inform you that this drug may cause some side effects. A first group of them includes nausea, which occurs infrequently and only in particularly sensitive patients, and headache which, again, occurs infrequently and only in particularly sensitive patients; these side effects are both serious. Then, you may have insomnia: it is not serious but can be frequent; however, once again I would like to reassure you that it occurs only in particularly sensitive patients. All these side effects can occur in a strong form.

5.2 Some Examples of Plan Revision Rules

Our attempt to solve the text revision problems outlined in Sect. 4 is far from being a general purpose one: more study is needed to better understand how to *read* a text plan in order to discover its argumentative deficiencies. For the time being, we have defined some *ad hoc*, strictly domain-dependent rules whose right sides include the following tree transformations: introduction of new subtrees or new intermediate nodes, change of the relative order of subtrees or change of the RR associated to a node. For example:

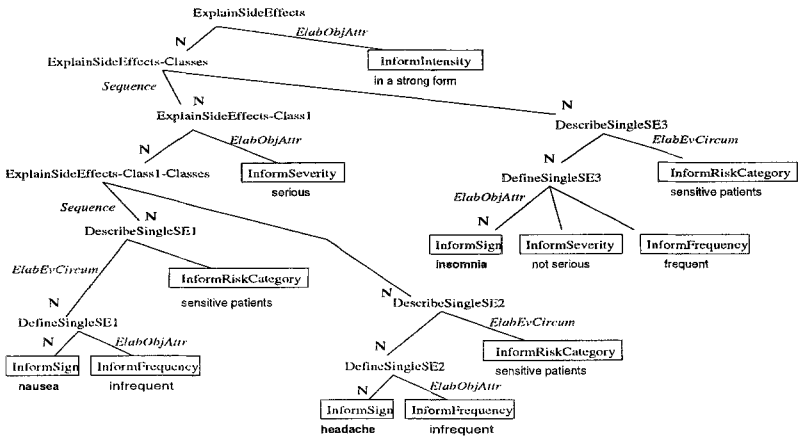

Fig. 4. The subplan tree of Fig. 3 after the treatment of repetitions

R3 IF the leaves of the subtree 'Persuade about the need of treatment' include several favourable items AND the patient has a 'hostile' attitude towards drug therapy THEN substitute the RR of Solutionhood with a Concession, in the main root of the tree (Ex1)

R4 IF at least one leaf in the subtree 'Request to administer the drug' is a favourable item AND the patient has a 'hostile' attitude towards drug therapy THEN substitute the 'Request to perform treatment' with a new subtree, to emphasise these items. (Ex3,4,5)

R5 IF at least one leaf in the subtree 'Enable to administer the drug correctly' is an item relevant for correct treatment AND the patient is aged THEN introduce a new subtree before the Enable, to emphasise these items. (Ex8,9,10,11)

6 Conclusions

We share Hovy and Wanner's opinion that "the generation process is even more complex than originally viewed, and that this complexity centrally affects the planning procedures" [8]. The tree-rewriting rules that we propose in this paper are domain-dependent heuristics, derived from analysis of empirical data. A domain-independent improvement in the discourse generation process could only be obtained by reconsidering Text Planning as a conjunctive and non serially decomposable problem, in which strong interactions between subgoals exist [9]. This would require representing *communicative actions* in a more fine-grained way than at present, by specifying their negative as well as their positive effects. In argumentative texts, in particular, graduality in the persuasion process could be obtained by introducing degrees of goal adoption and degrees of belief holding in the mental model of the addressee. Let us take, for example, the case of an 'Inform about side effects', as a leaf of the subtree of 'Persuade that the drug is useful and not harmful'. If the side effect is serious, this act will reduce

the addressee's intention to perform the treatment, that was increased by the previous subgoals of 'Persuade about the need of treatment' and 'Persuade that the drug is not harmful'. The combination of several, serious side effects will then produce a 'negative precondition' for the 'Enable to administer the drug correctly', which depends on their number.

With the Text Planning techniques that are applied at present, we claim that several plan revision tasks are needed to counteract negative effects of not considering subgoal interactions: sentence planning is only one of them. We acknowledge the importance of style refinement, but we contend that other, more argumentative tasks are important as well, and that knowledge about the domain and about the addressee plays a crucial role in all of them. In the medical field, this knowledge concerns not only what addressees presumably know, want or prefer, but also emotional factors such as, for instance, fear about disease prognosis or anxiety about pain.

References

1. D. Berry, I. Michas and F. de Rosis: Evaluating explanations about drug prescriptions. Submitted to Psychology and Health.
2. H. Dalianis and E. Hovy: On lexical aggregation and ordering. Proc. of 8th International Natural Language Generation Workshop (INLG-96), Poster session (1996) 29–32.
3. B. De Carolis, F. de Rosis, F. Grasso and A. Rossiello: Generating recipient-centered explanations about drug prescription. Artificial Intelligence in Medicine **8(2)** (1996) 123–145.
4. F. de Rosis, F. Grasso, C. Castelfranchi, I. Poggi: Modelling conflict resolution dialogues between believable agents. Working notes of the ECAI-96 Workshop on: Modelling conflicts in AI (1996).
5. D. Forsythe: Using ethnography in the design of an explanation system. Expert Systems and Applications **8(4)** (1995) 403–417.
6. I. Haimowitz: Modeling all dialogue system participants to generate empathetic responses. Computer Methods and Programs in Biomedicine **35** (1991) 321–330.
7. E. Hovy: Automated discourse generation using discourse structure relations. Artificial Intelligence **63** (1993) 341–385.
8. E. Hovy and L. Wanner: Managing sentence planning requirements. Working notes of the ECAI-96 Workshop on: Gaps and Bridges, New directions in planning and language generation (1996) 53–58.
9. R. Korf: Planning as search: a quantitative approach. Artificial Intelligence **33(1)** (1987) 65–88.
10. M. Meteer: Bridging the generation gap between text planning and linguistic realization. Computational Intelligence **7(4)** (1991) 296–304.
11. O. Rambow and T. Korelsky: Applied text generation. Proc. of 3rd Conference on Applied Natural Language Processing, Trento (1992) 40–47.
12. S. Reilly and J. Bates: Natural negotiation for believable agents. School of Computer Science, Carnegie Mellon University, CMU-CS-95-164, June 1995.
13. L. Wanner and E. Hovy: The HealthDoc Sentence Planner, Proc. of 8th International Natural Language Generation Workshop (INLG-96) (1996) 1–10.