Affective Natural Language Generation

Fiorella de Rosis[1] and Floriana Grasso[2]

[1] Dipartimento di Informatica - Università di Bari, Italy
derosis@di.uniba.it
[2] Department of Computer Science - University of Liverpool, UK
floriana@csc.liv.ac.uk

1 Introduction

The automatic generation of natural language messages (or NLG) has been employed in many computer systems and for various purposes, which go from providing information on a particular subject, to instructing on how to perform some -complex- action, to arguing about a particular claim (for an overview, see [24]).

In the majority of systems that have been designed in the last decade, these messages are adapted, in their content and presentation style, to the context in which they have to be applied: that is, to the Hearer's characteristics (represented in a "mental model" of the Hearer), to the application domain and, generally in a more implicit way, to the Speaker's characteristics. Adaptation is based on "strong" assumptions about the Hearer's mental state and the way this state is influenced by communication of each individual component of the message, and by understanding of the relationships among these components. The aspects of the mental state that are represented, in the large majority of cases, are the Hearers' beliefs and knowledge of domain topics and their goals about the domain state; in some cases, this may extend to representing other aspects, such as the ability to perform domain actions, the interest towards topics, the preference towards domain states. When a Speaker's model is represented as well, this includes second-order beliefs and goals of the same type.

Generally seen as informative tools, these systems give little space to representation of less rational aspects, such as emotions, of both Speaker and Hearer. However, natural language communication is influenced by a number of factors, of which the more rational ones constitute only a subset. Especially when communication occurs in "delicate" scenarios, such as a medical setting, the Hearer cannot be expected to coolly react to what is being said. Many studies in the health behaviour research have shown that patients' attention and understanding is highly affected by their emotional involvement [1, 26]. It therefore appears worthwhile to investigate whether, when and how emotions, personalities and other extra-rational factors should be taken into account when designing a NLG tool.

By paraphrasing Rose Picard [22], we define *Affective Natural Language Generation* as "NLG that relates to, arises from or deliberately influences emotions or other non-strictly rational aspects of the Hearer": these aspects (denoted with the generic term of "attitudes") include personality traits, emotions and highly-placed values. We argue for the need for affective NLG by considering an example from medical explanation texts, in which affect plays a particularly relevant role. We claim that Affective NLG requires a revision of the generation methods and of the knowledge sources (about the domain and about the Hearer) that are employed in this process, and, from analysis of the presented example and of the literature, we draw some conclusions about the way NLG methods might be revised to produce more affective texts.

2 A Case Study

We consider two examples from a corpus of explanations about drug prescription [2]. The corpus was collected by presenting a set of scenarios to several doctors, who were then asked to make a drug prescription and to support this prescription with an explanation addressed to the patient, to a nurse and to a colleague.

The two examples presented here are both addressed to the patient, as these were the versions of the explanation conveying more affect, and refer to the following scenario:

> *"A 44 years-old alcoholic man lives on a poor diet in a cold and damp lodging house. He now complains of a persistent cough, occasionally bringing up blood. His appetite is poor, and he thinks he is losing weight. Examination and chest X-rays, followed by other tests, confirm the diagnosis of tuberculosis".*

The explanation texts provided by the two doctors are shown in Fig. 1. Doctor D1 (column 2) provides a shorter text, Doctor D2 (column 3) a longer one. The two texts do not really differ in the overall structure: they both start with illustrating the disease, to then anticipate that treatment will be long and cannot be avoided; after synthesizing the main aspects of treatment (number of drugs and intake modalities), the texts end with mentioning the main accompanying measures that the doctor is going to undertake.

Both texts employ some typical *affect-conveying* techniques (italicized text in the figure). However, it is the quantitative and qualitative difference in the way these techniques are employed which makes the two texts really different:

- D1's text is mainly focused on convincing the patient to *follow the treatment*, by relying on his sense of responsibility: treatment is long and necessary (4) and has to be done every day, for at least three months (7).
- D2's text combines the goal of *motivating the patient* with the goal of *reassuring him* about the perspectives of success of treatment. This text is longer, first of all because it includes many more justifications:
  1. of diagnosis, and the way it is linked to available evidence (1),
  2. of the reason why a combination of drugs is needed (6),
  3. of the reason why a notification of the disease is required (not from the legal but from the public health viewpoint!: 10), and
  4. of the circumstance that may justify a future hospitalization (11).

  At the same time, the patient is constantly reassured about the efficacy of the treatment prescribed (2, 4, 5, 6) and about the health service participation to monitor his health status (11).

The first effect of the difference in affect, in the two texts, is therefore *a difference of content*: more affect implies, in the considered example, more details[1]. The second effect is a *difference of form*: in D2's text, many more affect-conveying terms are used:

- verbs: *kill off* the TB bugs, in segment 6; *a close eye on you*, in segment 11;

[1] In other cases, more affect implies, on the contrary, *less* details, caused by the decision to omit, from the explanation, aspects that are difficult for the Hearer to accept or that might cause distress to the Hearer, as we will see later on.

1	Mr Smith, *unfortunately* you have an infection of the chest which is called tuberculosis;	Well Mr Smith, we've got the tests back now and I think we might have mentioned your losing weight and being a bit poorly, *but we were a little bit worried* that you did have TB, or tuberculosis infection. And the tests have really confirmed that worry, it looks as if, on your chest X-ray, you've got active TB, and your sputum is also positive.
2		So, *the good news* is that we do have tablets that are *very effective* for treating TB,
3		*you do have to take several* tablets a day,
4	the problem, with this infection, is it *takes a very long time* to eradicate it from the body and therefore we have to undertake *quite a long course of treatment* which *it is essential* for you to fulfill for the full course.	and you will have to take them for some months *to get really over* this problem,
5		*but it is something we can do something about*.
6	Therefore, initially we are going to start off with 3, and *we may even* add 4 drugs when we talk to the microbiologist.	What we are going to do today is start you on two types of tablets, one a special antibiotic to *kill off* the TB bugs, the other ones are vitamins, because you can run *a bit short* of vitamins when you're on these tablets.
7	These drugs will have to be taken every day, and then we'll reassess the situation in three months time.	
8	I notice that you live on your own at the moment and	Now the other thing we've got to do is to look into why you've got this, *we're a little bit worried* that you're *maybe a bit undernourished*, and haven't been looking after yourself, and
9		we're going to see if we can speak to our social workers to see if they can improve on your surroundings and benefits.
10		What we're also going to have to do is notify you as a case of TB, because this is an infection and you might have caught it from one of your friends, or alternatively you could maybe cough on somebody and infect them.
11	I just wondered what you felt of coming into hospital over the next few days.	So, these are the main things we're going to be doing, and *we'll be keeping a close eye on you here at the clinic* to make sure you are getting better, and *we may need to bring you into hospital for a bit* if you don't progress as well *as we'd like*.

Fig. 1. Two explanation texts about drug prescription.

- adjectives: the *good news*, in segment 2;
- adverbs: a lot of *little bit* and some *very, really, maybe,* here and there, when needed to enhance positive aspects and mitigate negative ones.

Both texts (and especially D2's) are rather *redundant*: some topics are repeated with identical or equivalent wording. For instance:

- "it takes a very long time to eradicate", and "we have to undertake quite a long course of treatment" (D1's text, segment 4);
- "The good news is that we do have tablets that are very effective for treating TB", and "but it is something we can do something about" (D2's text, segments 2 and 5).

Other topics, on the contrary, are only touched on, without going into details: we will see some examples of this when we will talk about side effects, in particular.

Finally, both doctors employ the "first plural person" form of verbs, to give the patient the feeling that caring for the diagnosed disease is not something that concerns the patient alone, but it is something the doctor and the patient are managing together (in the case of D2, it is a "battle" they will fight jointly).

To conclude: affect manifests itself, in the examined example, both at the content and at the form level. Redundancy, inclusion of motivating and reassuring details, and elusion of demotivating topics are examples of the first level; use of enhancing or mitigating terms and of the first plural person form are examples of the second one. In other texts in the same corpus, we found several examples of a wise use of redundancy vs conciseness, to stress favourable or de-emphasize unfavourable information. We also found examples of summaries of topics that are relevant for correct treatment, that were introduced in the texts with the aim of reinforcing the patient's memory. Finally, we found a variation in the order of presentation of information items (for instance, side effects and drug administration) that was probably linked to their presumed impact on the patient's emotional state. In other examples of medical texts, other forms for expressing affect were employed: for instance, in a corpus of dialogues between doctor and patient in two different settings (a family planning service in Italy and an adolescent diabetic care service in the UK), several forms of *omission* were employed like elusion or deception [3].

All the examples we looked at, anyway, show that affect plays a crucial role in explanatory texts. We shall reflect, in the remaining of this paper, on how such a trait may be rendered in automated NLG.

3 Affect in Natural Language Generation

Apart from works which explicitly deal with (especially informal) argumentation [25, 27, 15, 10], reflection on how texts may be generated, that give appropriate weight to affective components of both the Speaker's and the Hearer's mental state, has been rather episodic in the NLG community, and mainly with a focus on the wording of the text.

Among others, [12] focuses on the idea of "producing utterances empathetic to both the Speaker and the Hearer", by offsetting "unpleasant" information and stressing "favourable" one, through detensifier and intensifier adverbs: the Hearer's mental model is enriched, to this purpose, with domain-related personal preferences, concerns, worries and related features. In [30] it is discussed how variables like the "social distance" between the Speaker and the Hearer, the "power" that the Hearer has over the Speaker and the two agents' "basic desires" might affect strategies for realizing a particular speech act. Similar considerations are made in other projects whose purpose is to build "life-like characters with personality and emotions", as, for instance, in [18].

Very few works have dealt with the problem of organizing the text content when affect has to be taken into account. In his book on *Generating natural language under pragmatic constraints*, Hovy dedicates a Chapter to this subject by discussing, in particular, how to deal with situations in which there is a "potential conflict" or an "agreement" between the affective implications of the text and the Hearer's opinion [13, Chapter 4]. He suggests to apply content-related and form-related techniques, the first ones aimed essentially at avoiding "sensitive issues" or at de-emphasizing them through evasion, selectivity or other means, the second ones aimed at enhancing positive aspects and mitigating negative ones, through appropriate word choices. Hovy's proposal to intervene on the text content through some form of evasion, elusion or "deception by masking" will remain an exception for some time. Kölln shifts this research field's focus towards the planning step, by suggesting to consider some Hearer's attitudes such as "subjective preferences" towards relevant concepts in the domain under consideration, in deciding the text content in concept explanation generation [16].

Some other researchers have been concerned with how emotional factors affect the structure of the sentence. One of the most notable contributions is Marilyn Walker's study on the role of repetitions in texts [28, 29]: this study goes in the opposite direction of the prevalent tendency in sentence planning, whose main focus is on *aggregation* [14]. In a recent survey, Reape and Mellish [23] report several definitions of this NLG task found in the literature, for instance: "redundancy elimination, abbreviation, text structure combination for concise and coherent text generation" etc. These goals may be achieved by language-dependent or language-independent techniques and may focus on conceptual, semantic or rhetorical aspects of the text. There are no considerations, in this paper like in others in the same field, on when aggregation should *not* be made. By illustrating the benefits of repetitions in a text, Marilyn Walker, on the contrary, indirectly indicates some circumstances in which redundancy is a value rather than a vice and therefore delimits the role of aggregation. We came to similar conclusions in our analysis of medical explanation texts and dialogues, in [6]: as the emotional state of the Hearers influences their reaction to the communicated message, assumptions that generally justify the avoidance of repetitions are likely to be violated in affective texts, and NLG methods should be revised accordingly.

Apart from these and few other examples, we lack, to our knowledge, a comprehensive research which fully accounts for how the goal of producing affective texts influences the various stages of the NLG process. In the next section we try to make a step in this direction, by sketching some ideas based on the literature and on our experience.

4 Considerations on How to Generate Affective Texts

If affect influences the information content and order and the way that communicative acts are realized, then all phases of the text generation process should be influenced by consideration of emotional factors. We assume a classic pipeline architecture for NLG systems, consisting of the phases of discourse planning, sentence planning and surface realization [24].

Fig. 2 illustrates this architecture as developed in our system: the Plan Library and the Augmented Transition Networks (ATN) library are knowledge sources employed in discourse planning and surface realization. A Speaker and a Hearer Model enable a double adaptation of the message produced; a Domain Knowledge Base (DKB) provides information on concepts, objects and other topics in the domain to which the explanation refers.

4.1 Discourse Planning

Discourse Planning is "the process of imposing ordering and structure over the set of messages to be conveyed" [24]. The result of this process, the text plan, is usually represented as a tree, whose leaves are the information content of the text and whose internal nodes have some representation of how parts of the text relate to one another.

Information content and order of the discourse items are often established by applying some planning algorithm to a Library of *plan operators*. The structure of the operators that have been applied in various discourse planning experiences varies slightly ([21] contains probably the most complete and clear introduction to this subject). We will refer, in this paper, to the structure of the operators that we employed in our tool for generating explanations about drug prescriptions [6]. In these operators:

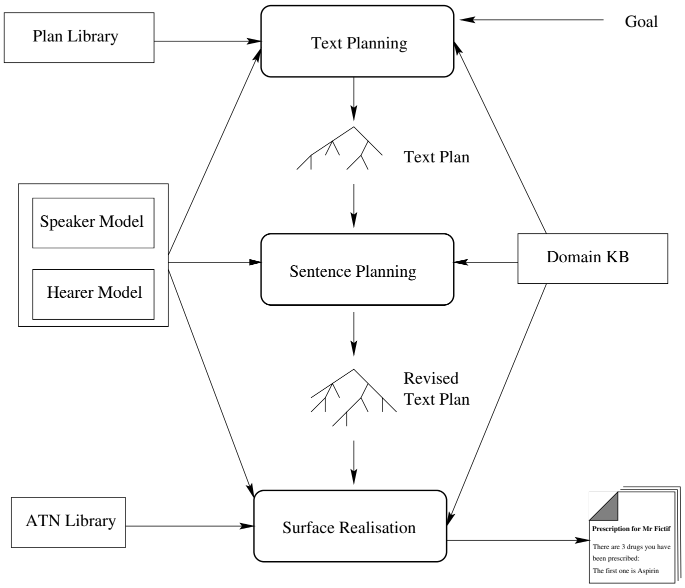

Fig. 2. Architecture for a NLG system.

- the *Header* specifies a discourse segment purpose;
- the *Domain Constraints* specify conditions on the Domain Knowledge base that must be satisfied for the operator to be applied;
- the *Intentional Constraints* specify conditions on the Hearer mental state that should be satisfied for the operator to be applied. As opposed to domain constraints, intentional constraints can in turn be planned for, if not satisfied, through a pattern-matching with the Effects of other operators;
- the *Effects* specify the changes produced in the mental state of the Hearer by the communicative action(s) activated by the operator;
- the *Decomposition* specifies how the communicative goal mentioned in the Header may be decomposed into sub-goals, and the role that each of these subgoals plays in the Rhetorical Relation (RR) attached to the operator[2];
- the *Rhetorical Relation* specifies the above mentioned RR name.

Figure 3 provides an example of one of these operators, that aims at describing a side effect (?x) of a drug (?y). The application constraint of this operator is that the Hearer (H) does not already know these side effects (NOT (KnowAbout H ?x)). Its effect is that the Hearer will come to know them. The goal specified in the Header may be decomposed into four subgoals (in this case, primary communicative actions) which are linked by a RR of *Elaboration Object-Attribute*; the goal to inform about the side effect name (Inform S H (Sign ?x)) is the nucleus of this RR; the other items in the Decomposition are satellites.

[2] We use the Rhetorical Structure Theory [19] to identify the relations among text segments: according to this theory, text segments participating to a rhetorical relation can have either a more prominent (nucleus) or a less prominent (satellite) role.

Header:	DefineSingleSideEffect (S H ?x ?y)	
Domain Constraints:	(Drug ?y) AND (SideEffect ?x ?y)	
Intentional Constraints:	NOT (KnowAbout H ?x)	
Effects:	(KnowAbout H ?x)	
Decomposition:	(Inform S H (Sign ?x))	*Nucleus*
	(Inform S H (Severity ?x))	*Satellite (optional)*
	(Inform S H (Frequency ?x))	*Satellite (optional)*
	(Inform S H (Intensity ?x)	*Satellite (optional)*
Rhetorical Relation:	ElaborationObjectAttribute	

Fig. 3. A typical plan operator, from [7].

In our generator of explanations about drug prescriptions, about 100 operators were employed to plan the message structure. Although this method enabled us to achieve a double adaptation of generated texts (to the Speaker and to the Hearer), these texts did not show the same level of empathy that we had found in the doctors' corpus of explanations.

To produce more affective texts, such a method should be revised throughout. First of all, the strategy to employ in achieving the communicative goal that is formalised in a plan operator has to consider not only the Hearers' experience or knowledge but also other, less rational factors, such as their desires, interests, fears, satisfaction and so on. Moreover, these factors should be 'graded' rather than being considered as dichotomic variables. This requires revising the Intentional Constraints and the Effects slots of plan operators.

Intentional Constraints. As said before, this is the slot in which the decision of when to include a topic in the message is formalised. In traditional text planners, this decision is driven by the need to avoid redundancy in the text, by avoiding to talk about what the user already knows. This originates from the Gricean maxim of "Quantity" [11] and is based on simplifying assumptions such as [29]:

1. unlimited working memory of the Hearer (everything an agent knows is always available for reasoning),
2. logical omniscience (every logical conclusion of the Hearer's knowledge is always available for reasoning) and
3. no autonomy in knowledge (assertions and proposals by the Speaker are accepted by default by the Hearer).

However, as stated by Marilyn Walker [29] and as we saw in our Examples in Sect. 2, it turns out that a topic may be included in a text, even if the Hearer is presumed to already know it, for several reasons: for instance, to augment the evidence in support of a desired belief, to make the topic salient in a given phase of the discourse, or to increase the Hearer acceptance of some claim. The need for mentioning a topic therefore depends, on one side, on characteristics of the Hearers that go beyond their knowledge state, such as their inference ability, their attentional capacity or their emotional state. But it also depends on characteristics of the discourse topic, such as how complex a task is, how hard and how important the inference is, or how acceptable errors consequent to misunderstanding of the message are. When the text has an argumentation purpose (for instance, "to convince the patient to carefully comply with treatment"), the target of the communication process are the Hearers' goals and intentions rather than, or in

Header:	DefineSingleSideEffect (S H ?x ?y)
Domain Constraints:	(Drug ?y) AND (SideEffect ?x ?y) AND (Eq ((Severity ?x) HIGH)
Intentional Constraints:	(Val(KnowAbout H ?x) $<$ t1) AND (Val(WantsToKnow H ?x) $>$ t2) AND (Val(Fear H PAIN) $<$ t3)
Main Effects:	Increases Val(KnowAbout H ?x)
Side Effects:	Decreases Val(Goal H (Take H ?y)) AND Increases Val(Fear H PAIN)
Decomposition:	(Inform S H (Sign ?x)) *Nucleus*
	(Inform S H (Severity ?x)) *Satellite (optional)*
	(Inform S H (Frequency ?x)) *Satellite (optional)*
	(Inform S H (Intensity ?x) *Satellite (optional)*
Rhetorical Relation:	ElaborationObjectAttribute

Fig. 4. A plan operator for *affective* discourse planning.

addition to, their beliefs and knowledge. The simplifying assumptions leading to avoid repetitions may be rephrased, in this case, as follows:

1. unlimited performance desire of the Hearer (everything the agent intends to do is immediately turned into action),
2. behavioural total coherence (every intention consequent of the Hearer's mental state is immediately adopted),
3. no autonomy in behaviour (requests are accepted by default by the Hearer).

To account for affective factors in the definition of plan operators' constraints, conditions under which these assumptions are likely to be violated, and attitudes of the Hearer's mental state that may affect them, need to be established.

Figure 4 shows an example of how the operator of Fig. 3 may be revised, to plan an affective discourse. This operator is meant to capture a situation in which the considered drug has "serious side effects" and the patient has a high " self-care attitude" and does not fear too much about painful consequences of taking the drug. New items are introduced in the Intentional Constraints slot: the Hearer's interest towards knowledge of health care implications (in our case, the side effects of the drug: (WantsToKnow H ?x)) and his/her emotional state. The particular emotion category considered in the example is 'fear' about 'the prospective undesirable event' of 'feeling pain due to the side effects of the drug' [9]. All conditions in this slot are transformed from first order formulae to conditions on the value of the 'intensity' of items: the operator will be activated only when this value (the 'level of knowledge', the 'level of interest', the 'level of fear') exceed or are below a threshold value. A different operator would consider the case of patients whose 'self-care attitude' is not very high or whose fear about feeling pain is high: in such case, details about most negative aspects of side effects will be hidden, or will eventually be attenuated by appropriate techniques.

Effects. The communication of a topic produces several effects on the Hearer's mental state: some of them are "main (usually positive) effects" that satisfy the main communicative goal of the Speaker, while others are "side effects" that may be "negative", not intentional and not avoidable. This makes too simplistic the on/off hypothesis, according to which agents either have or have not a goal, a belief or an attitude, and reverse their state of mind after receiving a communication. Some more fine grained representation is needed, instead, to express the graduality of mental state change.

Let us compare, again, the two operators in Fig. 3 and Fig. 4. In the operator employed in affective discourse planning, the main effect is to increase the Hearer's knowledge about the side effects of the drug. In addition, a new category of effects is introduced: the "Side Effects" slot (side effects of the communication process!) specifies that the consequences of being informed about the side effects (of drugs!) is a decreased intention of taking the drug (Goal H (Take H ?y)) and an increased fear about their possible negative consequences (Fear H PAIN).

4.2 Sentence Planning and Plan Revision

In the sentence planning step, a number of activities are performed, all having in common the characteristics of producing, from an input discourse plan, a new plan that is, in some sense, "optimized".

Aggregation, that is the grouping together of some text segments in order to enhance coherency and fluency, is considered the most important of these activities, in the sentence planners that have been proposed so far [5].

Highlighting of primary subjects and reordering of plan tree components are other 'affect enhancing techniques' that may be applied, if needed, after a plan has been produced and before surface realization starts.

These sentence planning algorithms should not, however, be driven only by style considerations and should only be applied in specific circumstances, that depend on the Hearer's affective state on one side, and on the domain topics on the other side.

Let us consider, for instance, aggregation. The decision of whether to aggregate sentences or clauses or whether, on the contrary, to emphasize repetitions in sentences, should consider not only aesthetic parameters, but also the *Hearer characteristics*, the *degree of difficulty* of the topic and its *importance*.

Example:
Let us suppose that the considered drug has three side effects (nausea, headache and insomnia), that each effect has its own severity, frequency and intensity and that the operator in Fig. 4 has been applied in a discourse plan in which side effects are described in sequence. A text with the following content will be produced in this case:

> *"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."*

This text includes several repetitions, some of which are concerned with "positive" aspects of this topic (*infrequently, in particularly sensitive patients*), other with "negative" ones (*it is serious, in a strong form*). To avoid increasing the patient fear about the consequences of taking the drug, a plan revision algorithm may be applied to enhance positive aspects through emphasized repetition and to mitigate negative ones through aggregation. The following revised text will be obtained:

> *"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."*

Similar considerations may be applied to the other sentence planning techniques we mentioned before.

4.3 Surface Realization

In the surface realization step, the way that a sentence is rendered (words choice, sentence structure and style employed) depends, as well, on the topic and on the Hearer characteristics. Affective text may be obtained by employing rule-based heuristics that define when and how empathy elements have to be introduced in the text. Let us see an example of these heuristics, from the previous text:

> **Example:**
> To produce the *"however, once again I would like to reassure you that...only..."* from the revised text plan, a rule of the following type is used:
> "IF the subtree Explain negative effects of a drug includes "relevant for compliance" items which take unfavourable values AND the patient is particularly fearful about this type of effects, THEN de-emphasize these items"

Similar rules would produce the various *little bit, very, really, may be, kill of* detensifier adjectives, adverbs and verbs in the example in Section 2.

5 From Theory to Practice

We have applied some of the principles outlined in the previous Section to our tool for generating explanations about drug prescriptions [6], with the aim of producing more affective messages. We did not revise, so far, the discourse planning component: this would require a theory for stating how *intensity* of knowledge, intentions and emotions should be updated after each operator has been applied, or, in other words, a definition of the Decrease and Increase functions in the two Effect slot of the plan operator in Fig. 4, which is beyond the scope of this work. We rather focused our efforts on the sentence planning phase, by specifically looking at the three affect-conveying operations encountered in our corpus: treatment of repetitions, highlighting of relevant subjects and reordering of text spans. To this aim, we introduced a plan revision step in the architecture of our generator, after extending the Speaker/Hearer Modelling components and the Domain Knowledge Base (DKB).

The Hearer Model is enriched with new features that formalise interests, preferences and a few emotion types (fear, hope, admiration, reproach and gratitude). Information items in the DKB are labelled according to hypotheses about relationships between the Hearer's knowledge and those aspects of the Hearer's mental state that the Speaker (the doctor) intends to achieve. For instance:

– an item ?x whose knowledge may affect the patient's *capability* to apply correctly the suggested treatment plan ?y is labelled as 'relevant-for-correct-treatment'. The cognitive hypothesis, in this case, is that:
(KnowAbout H ?x) $\rightarrow$ Increases Val(KnowHow H (Take H ?y)).
An example of such an item is 'drug intake intervals'.

- an item ?x whose knowledge may affect -either positively or negatively- the patient's *intention* to follow the suggested treatment, for instance to take drug ?y, is labelled as 'relevant-for-compliance' and may take 'favourable' or 'unfavourable' values. The cognitive hypothesis, in this case, is that for favourable items:
  (KnowAbout H ?x) $\rightarrow$ Increases Val(Goal H (Take H ?y))
  and that for unfavourable items:
  (KnowAbout H ?x) $\rightarrow$ Decreases Val(Goal H (Take H ?y)).
  Relevant for compliance items are related to *events* that may affect goal-based emotion categories: for instance, a 'serious' or 'frequent' side effect of a drug will increase the risk of 'feeling pain', which raises the Hearer's emotion of 'fear'. They may be related, as well, to *acts* that affect belief-based emotion categories or compound emotions: for instance, believing that the doctor 'cares about the patient's health state' may affect the patient 'gratitude' towards the doctor [9].

In our affective natural language generator, plan revision algorithms are activated by rule-based heuristics that specify when and how affect has to be added in a text. Examples of rules aimed at, respectively, treating repetitions, highlighting relevant subjects and reordering text spans are the following:

R1: IF the subtree 'Explain negative effects of a drug' includes unfavourable items with the same value, AND the patient is fearful about pain, THEN aggregate these repetitions.
R2: IF at least one leaf of the plan tree is a 'relevant-for-correct-treatment' item AND the patient is aged THEN substitute the 'Request to perform treatment' with a sentence which introduces the most significant of them.
R3: IF the side effects mentioned in the plan tree are many AND most of them are serious, THEN make sure that the administration details are mentioned before the side effects.

These rules are represented internally by an algorithm that explores the plan-tree in a top-down way, and finds out, in order, whether reordering, highlighting or treating repetition methods have to be applied.

The activated *plan revision* algorithms apply to the rhetorical structure of the plan tree to produce a new plan tree in which the intended operation is performed. These algorithms exploit several properties of the discourse tree, the main of which is the concept of "most nuclear part" defined by the Rhetorical Structure Theory [19], to guarantee that discourse coherence is preserved after plan revision:

- The module for *treating repetitions* classifies the leaves in the subtree of interest according to their role (nucleus or satellite) and to their value of 'relevance-for-compliance': a new subtree is then produced, in which topics with 'unfavourable' characteristics are grouped, repetitions of these characteristics are pruned out and the correct nuclearity or the discourse is restored.
- The module for *highlighting relevant subjects*, on the contrary, identifies 'relevant-for-correct-treatment' topics among the plan tree leaves, in order to produce the minimal subtree that generates a summary sentence in which these topics are mentioned. Such a subtree is subsequently grafted into the original plan tree, in an appropriate position.

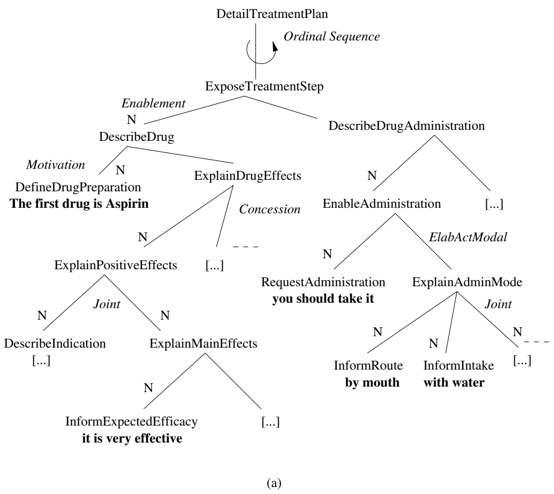

Fig. 5. Extraction of an highlighting subtree.

– Finally, the module for text spans' reordering applies properties of exchangeability of satellites in mononuclear RRs and of nuclei in some multinuclear RRs, to exchange the relative position of two text spans, by again insuring a final, coherent tree.

Figure 5 shows an example of highlighting subtree (on the right) that was extracted from the the part of the discourse plan that instructs on how to take a specific drug (on the left). The 'relevant-for-correct-treatment' topics are, in this example, the drug administration modalities (*by mouth, with water*); a 'relevant-for-compliance' item (drug efficacy) is also included in the minimal subtree, to increase the patient's intention to take this drug (as in segment 2 of Doctor 2's text, in Fig. 1).

The plan revision algorithms, the hypotheses on which they are based and the experiments which led to their design are described in detail in [6]. Extension of the affect-conveying techniques to the other phases of natural language generation (planning and surface realization) is a work still in progress.

6 Conclusions

In this paper we argued for the need for more research about how affect can be introduced in the automated natural language generation of explanatory and advisory texts, and we sketched some ideas that may be useful in this respect.

Far from advocating a flouting of Grice's maxims about natural language communication [11], we believe that Affective Text Generation requires interpreting Gricean maxims in a "relaxed" and extended way. When Grice suggests to "make your contribution as informative as required", this should not be interpreted as "never repeat something the user already knows, desires or may infer", but as "avoid repeating more than what is really needed". When he recommends to avoid "ambiguity" or "obscurity of expression", this should not be interpreted as "be always fully sincere", but as "avoid insincerity while not needed",... and so on. In this perspective, Grice maxims may be seen as default rules, and affective texts as exceptions in their application.

As a final observation, our discussion concerned how the user's emotions can influence a computer system producing text. A completely different issue concerns whether the computer system should show *its own* "emotions" in generating natural language. Studies have proven that, although conscious of being interacting with a machine, humans have a clear perception of the computer system's *personality* [20, 8]. On the other hand, it is still to be established whether users would appreciate a computer system showing the same level of empathy a human being would show, and in fact studies produced some evidence that this may not always be the case [4].

References

1. H.M. Becker, editor. *The Health Belief Model and Personal Health Behavior*. Thorofare, N.J: C.B. Slack, 1974.
2. D.C. Berry, T. Gillie, and S. Banbury. What Do Patients Want to Know: an Empirical Approach to Explanation Generation and Validation. *Expert Systems With Applications*, 8(4):419–428, 1995.
3. C. Castelfranchi, F. de Rosis, and F. Grasso. Deception and Suspect in Medical Interactions: Towards a simulation of believable dialogues. In Y. Wilks, editor, *Machine Conversations*, volume 511 of *International Series in Engineering and Computer Science*, chapter 8. Kluwer, 1999.
4. C. Cheepen and J. Monaghan. Naturalness in Automated Dialogues - less is more. In D. Levy and Y. Wilks, editors, *First International Workshop on Human-Computer Conversation*, pages 83–89, 1997.
5. H. Dalianis. *Concise Natural Language Generation from Formal Specifications*. PhD thesis, Royal Institute of Technology/Stockholm University, Department of Computer and System Science, June 1996.
6. F. de Rosis, F. Grasso, and D.C. Berry. Refining Instructional Text Generation After Evaluation. *Artificial Intelligence in Medicine*, 17(1):1–36, 1999.
7. F. de Rosis, F. Grasso, D.C. Berry, and T. Gillie. Mediating Hearer's and Speaker's Views in the Generation of Adaptive Explanations. *Expert Systems With Applications*, 8(4):429–443, 1995.
8. D.C. Dryer. Getting Personal with Computers: How to design personalities for agents. *Applied Artificial Intelligence*, 13(3):273–295, 1999.
9. C.D. Elliott, J. Rickel, and J. Lester. Lifelike Pedagogical Agents and Affective Computing: An Exploratory Synthesis. In M. Wooldridge and M. Veloso, editors, *Artificial Intelligence Today*, number 1600 in Lecture Notes in Computer Science, pages 195–212. Springer-Verlag, 1999.
10. F. Grasso, A. Cawsey, and R. Jones. Dialectical Argumentation to Solve Conflicts in Advice Giving: a case study in the promotion of healthy nutrition. *International Journal of Human-Computer Studies*, (to appear).
11. H.P. Grice. Logic and Conversation. In P. Cole and J.L. Morgan, editors, *Speech Acts*, volume 3 of *Syntax and Semantics*, pages 41–58. New York: Seminar Press, 1975.
12. I. Haimowitz. Modeling all Dialogue System Participants to generate Empathetic Responses. *Computer Methods and Programs in Biomedicine*, 35:321–330, 1991.
13. E. Hovy. *Generating Natural Language under Pragmatic Constraints*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1988.
14. E. Hovy and L. Wanner. Managing Sentence Planning Requirements. In K. Jokinen, M. Maybury, M. Zock, and I. Zukerman, editors, *Proceedings of the ECAI-96 Workshop on: Gaps and Bridges: New directions in Planning and NLG*, pages 53–58, 1996.

15. A. Jameson and T. Weis. How to Juggle Discourse Obligations. Technical Report 133, University of Saarbrüken, Department of Computer Science, October 1996.
16. M.E. Kölln. Employing User Attitudes in Text Planning. In *Proceedings of the 5th European Workshop on Natural Language Generation*, pages 163–179, May 1995.
17. J. Lewi and B. Hayes-Roth, editors. *Proceedings of the 1st International Conference on Autonomous Agents (Agent97)*. ACM Press, February 1997.
18. A. B. Loyall and J. Bates. Personality-Rich Believable Agents that Use Language. In Lewi and Hayes-Roth [17], pages 106–113.
19. W. Mann and S. Thompson. Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. *Text*, 8(3):243–281, 1988.
20. Y. Moon and C. Nass. How "Real" Are Computer Personalities? Psychological Responses to Personality Types in Human-Computer Interaction. *Communication Research*, 23(6):651–674, 1996.
21. J. Moore. *Participating in Explanatory Dialogues*. MIT Press, Cambridge (Mass.), 1995.
22. R.W. Picard. Does HAL Cry Digital Tears?: Emotion and Computers. In D. Stork, editor, *HAL's Legacy - 2001's Computer as Dream and Reality*, chapter 13, pages 279–303. MIT Press, Cambridge (Mass.), 1996.
23. M. Reape and C. Mellish. Just what *is* aggregation anyway? In *Proceedings of the 7th European Workshop on Natural Language Generation*, pages 20–29, 1999.
24. E. Reiter and R. Dale. Building Applied Natural Language Generation Systems. *Natural Language Engineering*, 3(1):57–87, 1997.
25. J.A.A. Sillince and R.H. Minors. What Makes a Strong Argument?: emotions, highly placed values, and role playing. *Communication and Cognition*, 24(3/4):281–298, 1991.
26. W. Stroebe and M. Stroebe. *Social Psychology and Health*. Mapping Social Psychology. Open University Press, Buckingham, 1995.
27. K. Sycara. Persuasive Argumentation in Negotiation. *Theory and Decision*, 28(3):203–242, 1990.
28. M.A. Walker. Redundancy in Collaborative Dialogue. In *Proceedings of the 15th International Conference on Computational Linguistics (COLING92)*, 1992.
29. M.A. Walker. The Effect of Resource Limits and Task Complexity on Collaborative Planning in Dialogue. *Artificial Intelligence*, 85(1-2):181–243, 1996.
30. M.A. Walker, J. Cahn, and S. Whittaker. Improvising Linguistic Style: Social and Affective Bases for Agent Personality. In Lewi and Hayes-Roth [17], pages 96–105.