PORTIA: A User-Adapted Persuasion System In the Healthy Eating Domain

Irene Mazzotta, Fiorella de Rosis and Valeria Carofiglio
Department of Informatics, University of Bari
Via Orabona 4, 70126 Bari, Italy
{mazzotta, derosis}@di.uniba.it

Abstract

Eating habits are influenced by emotional factors: persuasion to change wrong habits should therefore act on the central and the peripheral route at the same time, by combining rational and emotional strategies appropriately. In this paper, we describe a prototype system that is aimed at simulating user-adapted persuasion dialogues in this domain. We apply a theory of a-rational persuasion in which emotional and non emotional strategies are integrated, and describe how we collected and analyzed a corpus of messages in the domain and compared the persuasion strength of alternative strategies in this corpus. We then describe how we designed and implemented a prototype system which reasons on the interlocutor' mind to select the appropriate strategy to adopt in a given context, and translates the selected strategy into a natural language, rhetorically coherent text.

Keywords: user-centered design, natural language user interfaces, uncertain - probabilistic reasoning, health care applications

Corresponding author:
Fiorella de Rosis
Department of Informatics,
University of Bari
Via Orabona 4, 70126 Bari, Italy
http://www.di.uniba.it/intint/people/fior.html
derosis@di.uniba.it

1. Introduction

Eating habits are an essential component of wellbeing which result of cultural, psychological and educational factors. As such, they consolidate in time and, when wrong, are difficult to modify. Information media are masters in employing tricky arguments to persuade the population to consume products of doubtful healthiness. Attempting to persuade the population to adopt more appropriate habits by employing only 'rational' and 'scientific' arguments is probably not effective: this is one of the domains in which mingling of rational and emotional strategies is justified. Of course, as suggested by Walton (1992), attention should be paid to insure that arguments are relevant and strong: this is a subjective judgment which depends on the Persuader's knowledge about the Recipient and the conditions in which the message is uttered.

In this paper, we describe a prototype persuasion system that we called PORTIA[1] : this is the argumentation module of a dialogue system which acquires information about the users and then informs them about the advantages and disadvantages of various eating habits, provides suggestions in this domain and tries to persuade them to gradually modify problem behaviors. PORTIA employs three knowledge sources to select and formulate user-tailored persuasion attempts and argue about them: a *user model* including the presumed characteristics of the Recipient, a *persuasion knowledge base* to model rational and emotional strategies and *message plans* to represent how every strategy may be translated into a natural language message. Our theory of a-rational persuasion is briefly outlined in Section 2. The corpus of persuasion messages that was taken as a reference in the system design is described in Section 3. In Section 4, we outline the architecture of PORTIA and describe its knowledge sources and reasoning capability; we address the role of uncertainty in persuasion strategies and show how this is represented in belief networks and finally outline how the various strategies are rendered in natural language. Section 5 shows an example of system functioning, Section 6 adds some details about the system implementation and Section 7 concludes the paper by briefly pointing to open problem issues.

2. Background

In Walton (1990), the following statement by Govier is reported: *"Argument is a publicly expressed tool of persuasion. Typically it takes thinking to construct an argument. Reasoning is distinguished from arguing along these lines: reasoning is what you may do before you argue, and your argument expresses some of your (best) reasoning. But much reasoning is done before and outside the context of argument."*. This distinction between a phase of reasoning and a phase of formulation of the argumentation message is at the background of PORTIA. In the phase of reasoning, the system (that plays the role of the Persuader) works on a representation of the Recipient's mind to select a promising strategy, given its knowledge of the situation. In the phase of formulation of the argument, the selected strategy is translated into a discourse plan and then rendered with the media available (natural language or Embodied Conversational Agent).

In natural argumentation and persuasion, it is common to distinguish between 'rational' or 'cognitive' modes of persuasion and 'irrational' or 'emotional' ones. Miceli et al (2006) model how a persuader (P) builds a image of the recipient's (R) mind, in order to formulate a plan suited to influence this mental state by inducing the intention to perform a given action. In this model, the authors go over the mentioned distinction, to propose how 'rational' and 'emotional' modes of persuasion may be integrated to produce effective strategies in different contexts. By building on Toulmin's theory of argumentation and Walton's *argumentation schemes,* the authors identify two general modes of emotional persuasion: *persuasion through actual arousal of emotions* and *persuasion through appeal to expected emotions*. In the first case, the generative relation between emotions and goals is exploited. For instance, P's saying to R *'How disgustingly fat you are!'* is meant to provoke R's shame, which should generate R's goal of not losing his face and induce, as a means for this goal, his intention to go on a diet [2]. In the second case, the target is the Recipient's goal (not) to feel a certain emotion. For instance, P's saying to R *'If you are kind to Mary, you will feel at peace with your conscience (or you will not feel guilty)'* is meant to activate R's goal to feel at peace with his conscience (or not to feel guilty), in order to induce in R the intention to be kind to Mary as a means for it. In this theory, the *appeal to expected emotions* can be perfectly rational, as long as 'rational' implies the correct processing of the information available, the derivability of conclusions from premises and the production of plausible means-ends relationships. Therefore, an appeal to expected emotions is 'structurally' indistinguishable from any other 'argument from consequences', as defined by Walton (1992): the only difference resides in the *content* of the goal on which the persuader acts which, in this case, is that of 'feeling' a certain emotion rather than having a certain state of the world true.

3. Grounding The System On 'Natural' Data

In designing PORTIA, we wanted to integrate persuasion theories with observation of how humans behave when they wish to persuade someone to adopt a given behavior in the domain of healthy eating. To this aim, we performed two experimental studies: the first one was aimed at collecting a corpus of 'natural' data, while in the second one the perceived strengths of alternative strategies derived from that corpus were comparatively evaluated.

3.1. Corpus collection

In a web-based experiment, a scenario was presented to describe the situation in which the subjects involved (taking the role of Persuaders) should imagine to be. The scenario included the following hypotheses: friendship between Persuader and Recipient, Recipient's personality and goal, his living habits, his ability to make the action possible, relationship between desired action (eat vegetables) and likelihood to achieve the Recipient's goal. Two versions of the scenario were presented randomly to the subjects: one of them was formulated as a 'positive framing' (positive consequences of a diet rich in vegetables), the other one as a 'negative framing' (negative consequences of a diet poor in vegetables) (Levin et al, 1998). For more details about this study, see (Mazzotta and de Rosis, 2006): here, we will only synthetize the results which were useful for the system design.

We collected, overall, *32 messages* from Italian subjects with various backgrounds (psychologists, philosophers, computer scientists, epidemiologists, health care providers), aged between 23 and 63, of both genders and with different backgrounds. We factored every message into 'discourse segments': a segment could include one or more utterances with a given communicative goal. The average number of discourse segments per message (5.5) did not differ in the messages originating

---

[1] From the character of The Merchant of Venice, who was skilled in argumentation.

[2] In the whole paper, we will use a 'she' or a 'it' for the Persuader and a 'he' for the Receiver, without assuming any general rule about the distribution of roles between genders but only to simplify the presentation.

from positive and negative framing scenarios. To test whether framing affected the valence of arguments employed, we categorized every discourse segment as using ’negative’ or ‘positive’ arguments. The following are the main findings of our experiment:

- irrespectively of how the scenario was formulated, *the subjects tended to combine negative with positive arguments* but *preferred positive arguments to negative ones*: a large proportion of positive arguments (46 %) was employed also in the negatively framed scenario, while a lower proportion of negative arguments (26 %) was included in the positive framing case.
- Very few of the messages were formulated according to a ‘rational’ scheme: r*ational and a-rational arguments* were usually combined, with a *prevalence of emotional arguments* (57 %) both in the negative and in the positive framing conditions. Notice that the scenario was formulated so as to raise the subject’s attention on rational persuasion arguments (positive or negative effects on health of a diet respectively rich or poor of vegetables); we classified a discourse segment as ‘emotional’ when it included one of the techniques mentioned in (Miceli et al, 2006): ‘appeal to the goal to feel an emotion’, ‘emotional activation of a goal’, expression of emotion in the language style or display of some form of empathy. The following are examples of cases in which emotional arguments were employed:
    - in the *desirability of the goal* and its *activation*: “Just think Mary, how much more beautiful and healthy you would be if you ate more fruit and vegetables!” [3];
    - in the *proofs that conditions existed for making the activity*: “you are a really good cook!”
    - in introducing *high-order values* that the activity might contribute to achieve, such as ‘to live in a natural way’, ‘to satisfy gluttony’, ‘to enjoy’, ‘to make friends’: “You wouldn’t be hurting anyone and you’d be helping the biological farmers to live better” or “With vegetables you can prepare gorgeous, very light meals”;
    - in making more or less explicit *appeal to emotions*: “Here lies the wisdom of a mature woman: you have creative intelligence on your side *(pride)*, “If you insist on not eating fruit and vegetables, you show that you don’t care about yourself” *(self-esteem).*
- The *recommendation* of the behavior to follow was usually *introduced at the beginning of texts* which were prevalently rational, while it was introduced *only subsequently in more emotional ones*, after preparing the subject to receive the suggestion. In some cases, the role of this section became so minor, that it was not mentioned explicitly but was substituted with the description of some tempting consequences of the activity. This recommendation was supported with a *combination of different strategies*: by attempting to increase the desirability of the outcome, by reminding information about activity-outcome relationship or by proving that conditions hold to make the activity. E.g.: “A meal based on vegetables can be tasty: with just a little imagination you can prepare a first-rate dinner for your friends!”. Other segments were aimed specifically at evoking the *cognitive dissonance* in the receiver’s mind [4]. An example: “I’m surprised at you Mary! You… and then...”.

---

[3] All the examples in this Section are translated from Italian.

[4] According to the theory initially developed by Festinger, cognitive dissonance originates from inconsistency in the set of cognitive items held by a person: goals, beliefs, intentions, emotions etc. It is aroused whenever a person engages in an unpleasant activity to obtain some desirable outcomes and, being psychologically uncomfortable, motivates the person to reduce it (O’Keefe, 2002).

3.2. Comparative evaluation of strategies

To verify whether the persuasion strategies proposed by subjects in the previous experiment should be considered as a mere exercise of 'artifice production' or whether they could be seen as plausible and effective means to persuade, we performed a web-based evaluation study in which four strategies (presented in the form of a dialogue in English between two embodied agents) were compared. The study was a 2x2 (negative vs. positive framing and rational vs. emotional arguments) between subject design, with the same number of moves, the same duration and the same agents' expressions in the four conditions. Overall, 39 subjects from various countries were involved in the study: after displaying the dialogue, a questionnaire asked them to evaluate, with a Likert scale from 1 to 4, the agent's expression (*How much did you like the agent's performance?*) and the dialogue content (*If you were in the Receiver's shoes, would you be persuaded by the Persuader's words?*); two open questions enabled them to justify their evaluations. The following are the main results of this study:

- *the 'emotional and positive' version of the dialogue was considered as the most persuasive* on the average (2.4), the other three versions being equivalent (1.9 for the *'emotional and negative'* version, 2.0 for the two rational versions). The main critiques to the rational versions of the dialogue were that they were *"too much technical"*, that *"people don't talk like that, unless they are lecturing"*, that *"reasons employed were not enough strong"* or similar. The only critique to the 'emotional and positive' strategy was that it was *'too obvious'*, that *"Alice was too patently trying to convince John"*;
- the *'emotional and negative' version of the dialogue raised quite negative comments*: the scenery presented was seen as *'terrible'*, the Persuader was seen as *'violent'*, etc. The expected result was that the Recipient would become *'angry and defensive'*, and would stop listening;
- many subjects claimed that suggestions should be *"more tailored to the persuadee, less straightforward, more cautious"*, that the Persuader should have *"engaged the recipient in the discussion"*.

This evaluation study therefore proved the preference of non specialists in health promotion for a positive rather than a negative framing approach to persuasion. Consistently with the corpus analysis, it showed, as well, that purely rational argumentation was not seen as an effective method to persuade subjects in the domain of healthy eating, and that incorporating emotional issues was considered to be a more promising strategy. In all cases, adaptation of the message to the user characteristics was seen as a necessity. The two studies provided useful hints on how to design our persuasion system.

4. PORTIA

PORTIA is the persuasion agent of a domain-independent dialogue system: its architecture is outlined in figure 1. The agent includes a *user model* to reason on the presumed characteristics of the Recipient, a *reasoning module* to select and justify the appropriate persuasion strategy to a given Recipient and to repair it in case of failure, and a *planning module* to translate the selected strategy into a coherent language plan. We will now describe the various components in more detail.

FIGURE 1 ABOUT HERE

4.1. Reasoning on the Recipient mind to select a promising persuasion strategy

a. Persuasion strategies

In conditions of cognitive consistency, the main goal of a persuasion message is to *recommend the desired behavior* by influencing, at the same time, the attitudes which may affect positively the intention to conform to it (O'Keefe, 2002). This purpose may be achieved by acting on several aspects of the Recipient mind:

- by *influencing his values and goals,* be they 'rational'(e.g., to be in good health) or 'emotional' (e.g., anticipation of feelings such as *'Just think how satisfied you would feel after adopting a correct eating pattern.!'*),
- by *enhancing the perceived relevance of attitudes for behavioral changes* (*'You might not have realized it, but eating vegetables contributes considerably to improving your health')* or
- by *strengthening the Recipient's awareness of his ability to conform to the desired behavior* (*'You are a really good cook!'*).

In situations of inconsistency between the Recipient' attitudes and his behavior, he may be encouraged to a more consistent behavior by either inducing feelings of hypocrisy ("*You haven't been eating in accordance with your desire to be healthy, but now there is a chance to do so*") or by mentioning the (positive or negative) consequences of doing it ("*Here is a chance to act according to your attitude: and just think how bad you will feel if you don't take it"*).

Appeals matched with the Recipient's motivations are more likely to be successful than those engaging no salient desires. Knowledge of the Receiver's *wants* (preferences, goals, moral beliefs and significant values) is therefore essential in selecting the aspects on which to focus the persuasion process, that is, the outcomes the suggested behavior enables achieving.

b. Knowledge representation formalism

To Walton (1992), *"Practical reasoning is characteristically based on uncertainty or incomplete knowledge of a particular (changing) situation"*. We could verify, with our experiments, that representing uncertainty is essential when either emotional or rational strategies are applied. We therefore decided to represent persuasion strategies, as well as the Recipient's model, with belief networks (BNs); this is a well-known formalism to simulate probabilistic reasoning in directed acyclic graphs whose nodes represent random variables and whose oriented arcs represent any kind of relationship among variables (Pearl, 1988). A probability distribution is assigned to the variables associated with the 'root nodes' of the network (those which have no parents) and a conditional probability table to the other nodes.

c. The User model

Rather than acquiring this information about the Recipient (the system user) through direct questions, we attempt to implicitly infer them, with some level of uncertainty, from knowledge of his personality traits and living habits. The user model includes a *specific knowledge* and a *general knowledge* component. The specific knowledge collects facts about the user acquired during the dialogue. In the general knowledge, criteria to infer the user goals and abilities in conditions of uncertainty are represented in the form of elementary belief networks (EBNs), that is, networks with only one leaf node representing uncertain implications. In these networks, we represent evidence in support of the assumption that a given personality trait may be attributed to the Recipient, from a set of rules that we derived from personality assessment questionnaires [5]. Other rules represent the relationships between R's personality traits and his goals. The two sets of rules were built from the Big Five theory of personality traits and formalize the way personality factors (neuroticism, extraversion, agreeableness, conscientiousness, openness to experience) impact motivations of individuals, which in turn affect their performance.

Let us introduce the following notation:

- a is a variable denoting an action (e.g.: 'to eat vegetables'); g is a formula denoting a state of the world (e.g.: 'R is in good health');
- Bel, Int, A-Goal, V-Goal are modal operators that denote the various aspects of the mental state of R which are relevant in the persuasion process, that is, respectively, beliefs, intentions, active-goals and valued-goals. The first term of these operators denotes an agent, the second one is a formula. In particular:
- (V-Goal R g) stands for "g is a valued goal to R"; (A-Goal R g) for "R's goal g is active" [6]; (Bel R Implies(a,g)) for "R believes that doing a implies achieving g in a more or less near future"; (Bel R CanDo(R,a)) for "R believes that conditions hold for him to do a"; (Int R Do(R,a)) for "R intends to do a".
- The symbol '⇒?' denotes 'uncertain implication' and is represented in the BN with oriented arcs linking premises to conclusions. In the bayesian formalism, rule $(A1 \wedge A2 \wedge ... \wedge An) \Rightarrow ?\ B$ is interpreted as a conditional probability expression $P(B|A1,A2,..,An) = m$ stating that, among all the worlds satisfying A1 and A2 and...and An, those that also satify B constitute a fraction of size m. An uncertain implication specifies in a table the probabilities that B is true, conditional on all combinations of values for A1, A2, ..., An and enables assigning to the premises different weights in establishing the truth value of the consequence.

We may now make some examples of uncertain implications in the general knowledge component of our user model:

Extraverts enjoy being with people, are skilled in handling social situations, make friends easily etc:

Extraverted(R) ⇒?[EnjoyWithPeople(R)∧SkilledInSocialSituations(R)∧MakesFriendsEasily(R)]

Making friends is likely to be important to these subjects:

Extraverted(R) ⇒? (V-Goal R MakeFriends)

Agreable individuals respect others and have a good word for everyone:

Agreable(R) ⇒? [RespectOthers(R) ∧ HasAGoodWordForEveryone(R)]

Supporting farmers is likely to be important to these subjects:

---

[5] Myers-Briggs personality questionnaires are available at the following website: https://secure.wsa.u-net.com/www.teamtechnology.co.uk/tt/t-articl/mb-simpl.htm

[6] We basically share Cohen and Levesque's (1990) view of intentions in terms of goals the agent chooses and is committed to pursue. In addition, we introduce other properties of goals such as their being active or inactive: *"A goal is active when it is included in the agent's 'goal balance'; that is, when the agent starts to assess its importance and/or feasibility through comparison with other candidate goals, in*

Agreable(R) $\Rightarrow$? (V-Goal R SupportFarmers)

Goals can be inferred as well, from knowledge of the Recipient's habits. An example:

[MakeSport(R) $\wedge$ MakeCheckUps(R) $\wedge$ LookAtTv(R)] $\Rightarrow$? (V-Goal R GoodHealth)

Indidivuals who make sport regularly, undergo regular check-ups and are interested in medical TV programs are probably interested in being in good health.

Similar criteria are applied to infer whether conditions hold for the Recipient to perform the suggested action. An example:

[HasFreeTime(R) $\wedge$ LikesCooking(R) $\wedge$ AvailableVegs(R)] $\Rightarrow$? (Bel R CanDo(R,EatVeg))

Individuals who have some time free during the day, like cooking and live in a place in which good vegetables are available are probably in the condition to eat vegetables.

d. The persuasion knowledge base

Fragments of persuasion strategies are represented, as well, with elementary belief networks: these are classified according to the type of leaf-node:

- the generic *induction of intentions*-EBN represents the following relation:

[(V-Goal R g) $\wedge$ (A-Goal R g) $\wedge$ (Bel R Implies (a,g)) $\wedge$ (Bel R CanDo(R,a))] $\Rightarrow$? (Int R Do(R,a))

which may be read as *"if R has goal g and this goal is active, and R believes that doing a implies achieving g in a more or less near future, and R believes that conditions hold for him to do a, then probably R intends to do a"*.

This implication may be instantiated into several EBNs, each with an action which depends on the application domain and different (rational or emotional) goal; for instance: 'to be in good health', 'to make friends', 'to help the biological farmers to live better', 'to feel in good mood';

- the *induction of beliefs*-EBNs represent justifications of a belief, for instance about a means-end implication.

For instance, Argumentation from Examples may be rendered with an implication of the type:

[(Bel R Implies(a,g1)) $\wedge$ (Bel R Implies (a,g2)) $\wedge$...$\wedge$ (Bel R Implies (a,gn))]$\Rightarrow$? (Bel R Implies(a,g)), where g1,… gn are subgoals of g.

Arguing about (Bel R Implies(a,g)) can be represented, as well, in terms of Appeal to Expert Opinion, to Popular Opinion, or other.

Induction of beliefs-EBNs may be instantiated differently, by attributing a combination of values to the variables they include. An example:

[(Bel R Implies(EatVeg,StringMuscle&Bones)) $\wedge$ (Bel R Implies (EatVeg,TonifiedBody)) $\wedge$ (Bel R Implies (EatVeg,BetterSkin&Hair))]$\Rightarrow$? (Bel R Implies(EatVeg,GoodHealth))

may be read as "*if R believes that eating vegetables strengthens muscles and bones, tones up and rehydrates the body and is of benefit for hair and skin, then he probably believes, as well, that eating vegetables contributes to being in good health*".

---

view of its possible translation into an intention. An active goal may become an intention if that goal is finally chosen for pursuit." (Miceli et al, 2006).

e. Selecting a persuasion strategy

To simulate the reasoning followed by PORTIA in selecting the most promising persuasion strategy to a given Recipient, a complex BN is built by chaining forward several EBNs, starting from the data available about the Recipient. In this algorithm, EBNs are employed in a 'what-if' mode to test the persuasion strength of alternative candidate strategies:

- as a first step, the available evidence about R's characteristics is propagated in the EBNs whose child-node is a V-Goal (V-Goal-EBNs); the effect of propagating this evidence on the probability of every V-Goal node is observed: this probability is taken as an index of the 'degree of importance' of the associated goal, to R. Two stacks, registering separately the degree of importance (posterior probability) of rational and emotional goals are built. The V-Goal node g_i with the highest degree of importance is taken as the candidate goal on which to focus the persuasion strategy, and its EBN is selected from the KB;
- in the second step, a complex BN is built: the selected V-Goal EBN and all EBNs including R's characteristics and g_i are chained forward. Evidence is introduced in this complex BN, and the effect of propagating this evidence on the probability of the 'Int' node is observed:
  - if this probability exceeds a given threshold, then this is taken as the most promising strategy to R; in this case, a purely rational or purely emotional strategy will be applied;
  - else, the algorithm is applied recursively to the first goal g_j in the other stack: hence, if g_i was a 'rational' goal, g_j will be an emotional one (and the inverse). The two complex BNs are assembled through their common (Bel R CanDo(R,a)) node: in this way, a complex strategy that combines a rational goal (g_i) with an emotional one (g_j) will be applied.

Should the selected strategy fail, the two stacks will be employed to repair the failure, by selecting the next promising candidate goal (either rational or emotional).

4.2. Rendering the selected strategy in a coherent NL message

a. Generic plans as xml files

As we said in Section 2, after reasoning on the Recipient's mind to select the attitudes on which to ground persuasion, an argument must be constructed to express the selected strategy [7]. Items to possibly include in the argument correspond to the variables associated with nodes in the selected EBNs. The way these items are combined in the message (order in which to present them and relationships among the various parts) is represented in *argumentation plans* which are built on two theoretical grounds: Walton's argumentation schemes and Rhetorical Structure Theory (Mann et al, 1989). Although the principle behind an argumentation plan reflects the theory behind an argumentation scheme, their structure is not the same: some of the scheme's critical questions become preconditions in our plans, as well as their premises. The 'rhetorical relations' (RRs) linking preconditions among themselves and to the conclusion depend on the type of plan. Argument plans are represented as xml files. These are tree structures whose root represents the plan name and the situation in which it

[7] In agreement with (Miceli et al, 2006), we distinguish between 'persuasion', which is aimed at influencing directly an intention, and 'argumentation' , which is aimed at influencing some attitude behind intention. We will name our plans according to this distinction.

applies, and whose leaves correspond to communicative acts; a RR is attached to the root node and the intermediate nodes, to describe the relationship among children nodes.

Two variants of *persuasion plans* are built from Walton’s Argument from Positive Consequences: in both cases, the conclusion is expressed with a ‘Suggest’ communicative act (Suggest S Do(u,a)); this act is connected to the reasons which justify it ((Inform S Implies(a,g)) and (Remind S Like(R,g))) by a RR of ‘Motivation’. This part of the plan is connected with the statement that conditions exist to perform the suggested action (Claim S CanDo(R,a)) by a RR of ‘Enablement’. In the ‘Direct’ form, the suggestion (<c_act type=”Suggest” term=”Do(R,a)”/>) is presented as first, while the reasons that justify it (<c_act type=”Inform” term=”Implies(a,g)”/> and <c_act type=”Remind” term=”Like(R,g)”/>) come immediately after: a Remind is used when information was provided by the user, an Inform for new data. The claim that R is in the condition to perform a (<c_act type=”Claim” term=”CanDo(R,a)”/>) is left as last:

```
=================================================================
<plan name=”DirectPersFromConsequences” action=”a” goal=”g”>
  <RR name=”Enablement”>
    <RR name=”Motivation”>
      <c_act type=”Suggest” term=”Do(R,a)”/>
      <RR name=”Joint”>
        <c_act type=”Inform” term=”Implies(a,g)”/>
        <c_act type=”Remind” term=”Like(R,g)”/>
      </RR>
    </RR>
  <c_act type=”Claim” term=”CanDo(R,a)”/>
  </RR>
</plan>
=================================================================
```

In the ‘indirect’ form, the plan includes the same components but the suggestion is introduced after the supporting reasons. As we said, we saw several such cases of ordering of items in our corpus, especially when emotional strategies were employed.

```
====================================================================
<plan name=”IndirectPersFromConsequences” action=”a” goal=”g”>
  <RR name=”Motivation”>
      <RR name=”Joint”>
        <c_act type=”Remind” term=”Like(R,g)”/>
        <c_act type=”Inform” term=”Implies(a,g)”/>
      </RR>
      <RR name=”Enablement”>
        <c_act type=”Suggest” term=”Do(R,a)”/>
      </RR>
    <c_act type=”Claim” term=”CanDo(R,a)”/>
  </RR>
</plan>
====================================================================
```

Other plans are employed to support premises of persuasion plans. In these ‘argumentation’ plans, the conclusion is a belief or a goal. Two examples: the *goal-strengthening plan* is built from the Argument from Evidence scheme, to demonstrate that the selected goal is of value to R by listing a set of features Pi(R):

```
===========================================================
<plan name=” ArgumentFromEvidence” action=”a” goal=”g”>
  <RR name=”Evidence”>
     <RR name=”Joint”>
       <c_act type=”Remind” term=”P1(R)”/>
       <c_act type=”Remind” term=”P2(R)”/>
       .....
       <c_act type=”Remind” term=”Pn(R)”/>
     </RR>
     <c_act type=”Claim” term=”Like(R,g)>
  </RR>
</plan>
===========================================================
```

Belief-induction plans may formalize several schemes: e.g., an ‘*Appeal to Expert Opinion*’, to support the action-goal implication:

```
===========================================================
<plan name=” AppealToExpertOpinion” action=”a” goal=”g”>
  <RR name=”Evidence”>
     <RR name=”Joint”>
       <c_act type=”Inform” term=”Say(e,Implies(a,g))”/>
       <c_act type=”Inform” term=”ExpertSource(e,a,g)”/>
       <c_act type=”Inform” term=”BelievableSource(e,a,g)”/>
     </RR>
     <c_act type=”Claim” term=”Implies(a,g)”/>
  </RR>
</plan>
===========================================================
```

The same implication may be supported with a plan representing the Appeal from Position to Know scheme; an emotional variant of this scheme may be employed as well. Let us consider the following sentence: “*Your friend Paul tried the benefits of these simple, healthy foods on himself*”. In this case, the Persuader appeals to the experience of a person who is in the position to know (because he had a personal experience in the domain) and also a Recipient’s friend; this particular situation introduces a variant of the ‘rational’ scheme (that we call *Appeal from a friend’s personal experience*) in which friendship adds some ‘emotional strength’ to persuasion. This schema is represented in PORTIA as follows:

```
=================================================================================
<plan name="AppealFromFriendlyPersonalExperience" person="x" action="a" goal="g">
  <RR name="Evidence">
     <RR name="Joint">
        <c_act type="Inform" term="Say(x,Implies(a,g))"/>
        <c_act type="Inform" term="FriendOf(x,R))"/>
        <c_act type="Inform" term="PersonalExperience(x,a,g)"/>
     </RR>
     <c_act type="Claim" term="Implies(a,g)"/>
  </RR>
</plan>
=================================================================================
```

b. Complex plans instantiation and building

When a given goal (for instance: ‘to be in good health or ‘to preserve the environment’) has been selected as the goal on which to focus the persuasion strategy, the variable g in the persuasion plan-file is instantiated with the corresponding value. If, for instance, the selected goal was ‘to be in good health’, the instantiated ‘direct’ plan produces the following message:
“You should try to increase the quantity of fruit and vegetables in your diet! They have proved to be very effective for one’s health, which you seem to value greatly. I’m sure you can do it if you try”.

If, on the contrary, the goal was ‘to preserve the environment’ and the selected form was ‘indirect’, the message would be something like:
“It’s clear, John, that you feel very involved in respecting the environment and promoting a natural way of life and you know for sure that a diet rich in vegetables means respecting animals and the environment much more than a diet rich in meat! So, just think how satisfied you would feel if you increased the quantity of fruit and vegetables in your diet! I’m sure you can do it if you try!”

In addition to this basic part of the message, the user may be reminded of how important it is, to him, the selected goal, by providing evidence about the aspects of his behavior which demonstrate this hypothesis. This may be obtained by integrating the persuasion plan with a goal-strengthening argumentation plan. For example: the following complex plan is obtained by substituting the `<c_act type=”Remind” term=”Like(R,g)”/>` in the DirectPersFromPositiveConsequences plan with an appropriately instantiated ArgumentFromConsequences plan.

```
==============================================================================
<plan name=”Persuade&JustifyAboutGoodHealth” action=”EatVeg” goal=”GoodHealth”
type=”Indirect”>
  <RR name=”Motivation”>
      <RR name=”Joint”>
   .....................................................
        <RR name=”Evidence”>
          <RR name=”Joint”>
              <c_act type=”Remind” term=”MakeSport(R)”/>
            <c_act type=”Remind” term=”MakeCheckUps(R)”/>
          </RR>
        <c_act type=”Claim” term=”Like(R,g)”/>
        </RR>
   .....................................................
        <c_act type=”Inform” term=”Implies(a,g)”/>
      </RR>
    <RR name=”Enablement”>
        <c_act type=”Suggest” term=”Do(R,a)”/>
        <c_act type=”Claim” term=”CanDo(R,a)”/>
    </RR>
  </RR>
</plan>
==============================================================================
```

Complex plans have not necessarily to be rendered entirely, but may be pruned out in their less important parts to avoid too complex messages. The following message is an example of applying this pruning procedure to the previous plan:
“You do sport and look after your health with regular medical check-ups. Why don’t you try to increase the quantity of fruit and vegetables in your diet? I’m sure you can do it if you want to”.

5. An example

To demonstrate the functioning of PORTIA (in particular, the way uncertainty is handled), let us consider the following case: John (the Receiver of the persuasion message in this example) declared, at some point in the dialogue, that he makes sport and periodical medical controls regularly, enjoys being with people and makes friends easily; at the same time, his diet seems to abound in meat and carbohydrates and to be poor in vegetables and fruits. PORTIA builds a belief network to represent the presumed John's mind (figure 2).

FIGURE 2 ABOUT HERE

It propagates in this model the available knowledge: it infers that John is probably quite extraverted and therefore likes making new friends; at the same time, being in good health is probably an important goal to him. The two goals (V-Goal John MakeFriends) and (V-Goal John GoodHealth) have high probability values but the second one is more likely (figure 3, leftside).

FIGURE 3 ABOUT HERE

The system may then select between a 'rational' persuasion strategy focused on the goal of 'being in good health' or an 'emotional' one, focused on the goal of 'making friends'. It finds that either strategy, although being promising, does not seem to induce in John the desired level of intention to adopt the target behavior (to increase his consumption of vegetables): it therefore combines the two strategies as shown in figure 2:

[(V-Goal John GoodHealth)∧(Bel John Implies(EatVeg,GoodHealth))∧(Bel John CanDo(John,EatVeg))]⇒?
  P1-(Int John Do(John,EatVeg)) [8]

[(V-Goal John MakeFriends)∧(Bel John Implies(EatVeg,MakeFriends))∧(Bel John CanDo(John,EatVeg))]⇒?
  P2-(Int John Do(John,EatVeg))

By reasoning in a 'what-if' mode, the system propagates in the BN the assumption that John knows about the benefits of eating vegetables and is in the conditions to adopt this eating behavior, and discovers that he should very likely intend to do it (right side of figure 3: p=.80). However, PORTIA knows that John does not eat much vegetables: he is then probably in a condition of inconsistency among his beliefs, goals and intentions.

PORTIA has still several decisions to make before formulating its persuasion message in natural language. It must decide:

- whether to apply an *appeal to positive or to negative consequences* argumentation scheme;
- whether to include some form of *encouragement to a more consistent behavior*;
- whether to *justify its suggestion* or not, that is whether to mention, in the message, the premises that motivated its choice: John's goals, the behavior-consequences relationship and John' presumed ability to eat vegetables;

- in which order to present the two arguments;
- how to formulate the argument in natural language.

PORTIA decides to present as first the rational argument, which corresponds to the most important goal to R, and to integrate it with an emotional component presented in less detail, to formulate them both in positive terms and to employ a ‘friendly’ language in the message. To reinforce John’s awareness of the positive effects of eating vegetables on health and sociability it decides, as well, to include these items in the message. The final result is the following:

“I’m surprised at you John! You do sport, look after yourself with regular medical check-ups, then you eat a lot of meat and carbohydrates, almost excluding vegetables from your diet! Perhaps you don’t know the benefits that a diet rich in vegetables can have on your health. A dinner of fresh, tasty salads is easy to prepare and is an excellent way of having a good time with your friends.”

To build this message, we first combined the two complex discourse plans: ‘Persuade&JustifyAboutGoodHealth’ (Section 4.2.b) and ‘Persuade&JustifyAboutMakeFriends’. We then pruned off the suggestion as well as the claim about the rational goals in the first plan, and evoked the ‘emotional’ part with only the means-end implication. Appeal to cognitive dissonance (O’Keefe, 2002) was introduced at the beginning of the message in the form of system’s ‘surprise’ about inconsistency between John’s general and dietary behavior:

Opening of appeal to cognitive dissonance:

I’m surprised John!

<c_act type=”Remind” term=”MakeSport(John)”/>

you do sport

<c_act type=”Remind” term=”MakeCheckUps(John)>

you look after yourself with regular check-ups,

<c_act type=”Claim” type=”Like(John,GoodHealth)”>: omitted

Closing of appeal to cognitive dissonance:

you eat a lot of meat and carbohydrates, almost excluding vegetables from your diet!

<c_act type=”Inform”> Implies(EatVeg,GoodHealth)>

Perhaps you don’t know the benefits that a diet rich in vegetables can have on your health.

<c_act type=”Suggest” term=”Do(John,EatVeg)”>: omitted

<c_act type=”Remind” term=”CanDo(John,EatVeg)”/>

A dinner with fresh, tasty salads is easy to prepare

<c_act type=”Remind” term=”Implies(EatVeg,MakeFriends)”/>

is an excellent way of spending a goodtime with your friends.”

6. Implementation

[8] The two nodes P1-(Int John Do(John,EatVeg)) and P2-(Int John Do(John,EatVeg)) represent the two (rational and emotional) strategies to persuade John to eat vegetables. These nodes are in a relationship of ‘nearly-and’ with (Int John Do(John,EatVeg)).

The *Persuasion Strategies KB* and the *User Model KB* are represented as sets of object-oriented belief networks enabled by Hugin[9] and manipulated with Hugin's APIs, which provide a powerful tool for dynamic building of BNs from EBNs; the *Argumentation Plans KB* are represented as sets of XML files and are manipulated with xml technology. The whole module is implemented in Java: this platform independence insure its reuse in other implementation contexts and dialogue systems.

As we said in Section 4, PORTIA is the persuasion component of a domain-independent advice-giving dialogue system. This system performs several tasks: it acquires information about the user, provides information on request or according to its own plans, suggests lines of action when appropriate, tries to persuade the user to follow them when needed and, finally, enters in an argumentation subdialogue to justify and support its choices or revise them if needed. The dialogue system's structure is based on the information state: a blackboard (the 'information state') plays the role of a common memory to all the modules; logical rules are employed to update this memory and select the next dialogue step. Other system components apply, on the contrary, probabilistic reasoning: PORTIA, as we said, but also the user modeling agent which, in addition to the tasks we considered in Section 4.1c, recognizes the 'social' attitude displayed by the user during the dialogue and adapts the system behavior accordingly. These requirements have been satisfied by developing the dialogue system with an agent-based architecture in JADE (Java Agent DEvelopment Framework) and by using CLIPS native methods by JNI (Java Native Interface) to simulate logical reasoning and Hugin's API to simulate uncertain reasoning in BNs. Representing message plans as xml files enhances domain-independence, modularity and flexibility of our system: complex and instantiated plans may be built from general and elementary ones by means of the xml technologies. A simple surface generation module translates communicative acts and rhetorical relations in the complex plans into natural language sentences, by introducing some adaptation in the selected style. The message may be labelled, as well, with an appropriate markup language, to be given as input to an Embodied Conversational Agent player implemented with Haptek [10] and Loquendo [11].

7. Conclusions

After outlining a theory of the possible interaction between emotional and non-emotional persuasion, in this paper we described how we designed a persuasion system in the domain of healthy eating, after collecting and analyzing a corpus of persuasion texts and evaluating their persuasion strength. We found that purely rational strategies very employed very infrequently while emotional elements could be found everywhere, in various forms; positively formulated messages were considered to be more persuasive, and adaptation of the message content and style to the Recipient an absolute need. Due to the high level of uncertainty in persuasion strategies, a probabilistic knowledge representation was adopted. Considering probability theory and belief networks as a method for treating uncertainty is not a novelty in the argumentation community: other authors proposed to measure the *strength of support* or the *confidence* the inference confers to an argument in probabilistic terms or applied belief networks to describe arguments in medical genetics. The readers may refer to (Reed, 1997), for an exaustive and clear survey of the main techniques for reasoning with argument, among which they describe Zukerman's NAG, the system who first employed BNs to represent argumentation strategies. In this paper, we demonstrated that this formalism enables representing forms of reasoning including various kinds of rational and emotional appeals. A

[9] www.hugin.com
[10] www.haptek.com

problem still open in PORTIA concerns the last step of message generation, in which the message plan is translated into natural language. So far, we adopted a simplified solution for this component, by using canned texts; although we made an effort in this direction, we could not render the variety and subtlety of language styles we found in our corpus: this merits a separate research project. We are aware of the limits of our work, due to including in our experiments laypeople rather than counselors or therapists and designing the system according to their knowledge. However, the final goal of our research is to build a conversational character which provides suggestions in this domain: we believe that such a character might hardly be seen as a substitute of the therapist. It might rather play the role of a 'competent friend', who knows about the addressee and exploits this knowledge to select a promising strategy to get the desired effect.

Acknowledgements

This work was financed, in part, by HUMAINE, the European Human-Machine Interaction Network on Emotion (EC Contract 507422). We acknowledge Giuseppe Clarizio and Vincenzo Silvestri for cooperating to develop the dialogue system and PORTIA.

References

1. D.N. Walton, *The place of emotion in argument,* University Park: Pennsylvania State University Press, 1992.
2. D.N. Walton, "What is reasoning? What is an argument?," *The Journal of Philosophy*, Aug. 1990, vol. 87, no. 8, pp. 399-419.
3. M. Miceli, F. de Rosis, and I. Poggi, "Emotional and non emotional persuasion," *Applied Artificial Intelligence: an International Journal*, Nov. 2006, vol. 20, no. 10, pp. 849-879.
4. I.P. Levin, S.L. Schneider, and G.J. Gaeth, "All frames are not created equal: a typology and critical analysis of framing effects," *Organizational Behaviour and Human Decision Processes*, Nov. 1998, vol. 76, pp. 149-188.
5. I. Mazzotta and F. de Rosis, *Artifices for persuading to improve eating habits*, AAAI Spring Symposium on "Argumentation for consumers of health care", tech report SS-06-01 No. ISBN 978-1-57735-262-4, Mar 2006, pp.76-85.
6. D.J. O'Keefe, *Persuasion: Theory and Research,* Sage Publications Inc., 2nd ed., 2002.
7. J. Pearl, *Probabilistic Reasoning in Expert Systems: Networks of Plausible Reasoning*, Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1988.
8. P.R. Cohen and H.J. Levesque, "Intention is choice with commitment," *Artificial Intelligence*, Mar. 1990, vol. 42, pp. 213-261.
9. W.C. Mann, C.M. Matthiesen and S.A. Thompson, *Rhetorical structure theory and text analysis*, Information Sciences Institute Research Report ISI/RR-89-242, Nov. 1989, pp. 89-242.
10. C. Reed, "Representing and applying knowledge for argumentation in a social context," *AI & Society*, Nov. 1997, vol. 11, 1-2, pp 138-154.

[11] www.loquendo.com