

From Language to Thought: Inferring Opinions And Beliefs From Verbal Behavior

Fiorella de Rosis and Nicole Novielli¹

Abstract. In this contribution, we describe our ongoing work in the direction of studying how negative or positive opinions may be recognized from language and how beliefs may be dynamically inferred from expressions of opinion. We begin by considering the language processing methods which have been applied to 'sentiment analysis' to show the results they produced and their limits, and then reflect on how beliefs may be inferred gradually, in conditions of uncertainty and by carefully considering various forms of context.

1. INTRODUCTION

Humans may express their opinions with several means: with actions, body attitudes and language: they may patently shiver, close the windows or say '*Cold today, isn't it?*', to manifest their opinion that the temperature at home is not adequate. Opinions may be about the environment (as in the previous example), about other people or about themselves. However, the relationship between beliefs and actions, attitudes and language is not so strict: I might simulate shiver, close the window or say '*Cold today isn't it?*' for reasons of politeness, because I presume that my partner living with me feels cold. Considerable efforts are being made towards inferring goals from observation of nonverbal behavior (see, e.g., [6]). Language is particularly difficult to interpret, as an expression medium: humans may more easily lie or simulate their beliefs by speaking than with their body expressions. And still, language will be, probably for years, one of the most common communication media with smart environments. In this short contribution, we describe our ongoing work on the problem of how negative or positive opinions may be recognized from language and how beliefs may be inferred from expression of opinions. We start from considering the language processing methods which are applied in 'sentiment analysis', to show the results they produced and their limits. We then reflect on the fact that beliefs may be built gradually, both in their strength and their level of certainty. We reason on the factors which may influence masking expressions of beliefs in various contexts and how these factors may be considered in the interpretation of a given sentence.

2. BACKGROUND

Relation between language and thought was the object of philosophers' and psychologists' research since long. Three phases may be recognized in the process of inferring the content of a speaker's thought by a hearer: a. identification of the *meaning* of words used; b. identification of the *proposition* expressed in light of the meaning and the rest of the situation in which the utterance

takes place and c. identification of further *implicatures* over and above the proposition expressed [5]. Phase b. is generally made possible only by the analysis of *context*, without which sentences would be not interpretable; the same is true for phase c., in which aspects of the mental state of the speaker, such as beliefs and intentions, are inferred. This inference goes beyond the immediate meaning of the utterance. In a well-known example by Grice [7], A may say to B '*the bus will be here within five minutes*' not just in order to transmit his belief, but in order to put her at ease, because he observed that she is impatient for the bus to arrive. But he might say it for a number of other reasons: to put himself at ease because he is impatient, as a pretext to begin a conversation with B, to justify the bus delay because he feels responsible for this, and so on. He might even say something he doesn't believe, either intentionally or without intentionally misrepresenting himself [8] and computers may imitate this behavior [1]. If B knows about A, for instance because the utterance to interpret was pronounced in the context of an ongoing dialogue, finding the most likely interpretation of A's utterance will be easier but the 'context' to consider will be wider. In defining a communication language among artificial agents, Cohen and Levesque [2] neatly stated the semantics of 'illocutionary acts' in terms of the effects the Speaker intends to achieve with them: the hypothesis was that this effect always consists in 'communicating own mental state', with the Speaker's 'sincerity' as a strong assumption about communication conditions. This work on communication language of artificial agents is of primary importance in the immediate interpretation of a given sentence in terms of an agent's beliefs and intentions; however, it is of more limited use in natural language understanding where (as we said in our previous examples) consideration of the context -in a wide sense- is essential to avoid trivial interpretations.

A rough description of the user's beliefs in a human-computer conversation could be made by just extracting and summarizing the opinions expressed during the dialog [1]. This simple summarization still requires considering the degree of uncertainty in the expression of opinions and of consistency in opinions expressed at different times. A more sophisticated description of beliefs requires, however, a wider consideration of the context in which the opinions were expressed and of other sources of knowledge about the speaker's mind.

3. OUR STUDY

Our long-term goal is to build a dialogue system which provides user-tailored suggestions about healthy living habits. According to Prochaska and Di Clemente's Transtheoretical Model of Change (TMC in [14]), this kind of dialogue should apply a strategy in which the presumed 'stage of change' of the client (from a 'wrong'

¹ Department of Informatics, University of Bari Via Orabona 4
70126 Bari, Italy - email: {derosis, novielli}@di.uniba.it

to a more ‘correct’ behavior) is considered, to adapt dynamically the information and persuasion plans to the specific situation. By building it in a domain-independent way, our ambition is to apply this dialogue system to any behavior problem occurring in a smart environment (smoking in presence of children, using energy sources in a not economic way etc).

The TMC model suggests how stages of change may be recognized from a set of signs which include the *value* given to the ‘correct’ behavior which is the object of the persuasion process, the *knowledge of reasons* influencing the adoption of a problem behavior, the *belief that* (internal and external) *conditions exist to change this behavior* and the *intention* to change it. Recognizing the stage of change in natural language dialogues therefore requires inferring these aspects of the users’ mental state from their verbal behavior. A corpus of dialogues may help to define some recognition methods and validate them. Our corpus of

conversations about healthy dieting with an ECA was collected with a Wizard of Oz tool [4]: the examples in this paper are extracted from that corpus. We will focus our reflections, in particular, on the expression of the *belief* that own behavior is wrong.

4. THE METHOD

To guess another person’s beliefs and intentions, one may ‘*perceive situational information about another and use it as input into one’s own cognitive structure*’ [6]. If the user’s beliefs to recognize regard his own behavior (whether it is ‘correct’ or ‘wrong’), perception may be translated into identification of expressions of opinions, while beliefs can be inferred through a mapping with the system’s belief structure.

Table 1. Excerpt of an example dialogue from our corpus

N	Adjacent pair	Recognition	Question
1	S: Do you think your diet is correct or would you like to change your dietary habits? U: I think I should control my consumption of sweets.	U declares that his <i>consumption of sweets is not correct</i> .	Does U really believe it?
2	S: Have you ever desired to change your diet? U: Yes	U declares he <i>wanted to change</i> his dietary behavior <i>in the past</i> .	Did U want it really?
3	S: Do you think your weight is correct or would you like to change it? U: I am 4 kg overweight so I would like to change it. What should I do?	U recognizes he <i>is overweight</i> . He manifests his (light) <i>intention of losing some weight</i> . He <i>asks some support</i> in this direction.	Does U really believe it? Has he really the intention to change? Does he really need support?
4	S: Why you never tried to do it? U: Because I’m not so patient	U declares he is not in the conditions to change his dietary behavior because of <i>lack of endurance</i> .	Does U really believe that he lacks of endurance?
5	S: Do you think you assume all the substances you need? U: I regularly assume carbohydrates, proteins and vitamins. I probably assume not enough minerals and too much fat.	U declares that his diet is: - <i>well balanced in carbohydrates, proteins and vitamins</i> ; - <i>not well balanced in minerals and fats</i>	Does U really believe that his diet is correct in these substances and incorrect in others?
6	S: You don’t have to completely eliminate fat foods. You have just to limit their proportions U: I can’t organize my diet because I have no time.	U declares that he is not in the conditions to change his dietary behavior because of <i>lack of time</i>	Does U really believe that he has no time to organize his diet?
7	S: Sure, organize your diet could appear a challenging task especially when nobody can help you and you have to work or to study. U: I am 90% invalid	U declares that he is not in the conditions to change his dietary behavior because of <i>inability</i>	Does U really believe that he is not able to prepare correct meals?

Let us consider the excerpt from one of the dialogues in our corpus, that is shown in Table² 1: we will go through this example (and some variants) to discuss some of the problems in progressively inferring the user’s opinion and beliefs as far as this short dialogue goes on.

4.1 Identifying expression of opinions about own behavior

Identification of positive or negative opinions expressed linguistically can be seen in terms of ‘sentiment analysis’. This method aims at recognizing the viewpoint underlying a text span:

a typical task is the binary classification of texts in order to define their polarity (positive vs. negative, that is ‘thumbs up’ or ‘thumbs down’, good or bad). This goal is achieved by applying traditional machine learning techniques to a multidimensional representation of the collection of documents. The definition of the set of features involved in the representation is crucial, and several groups are working on the selection and interpretation of indicators to improve results in terms of accuracy.

In the *bag of words* (BoW) approach, basic features for the vectorial representation are unigrams, bi-grams or tri-grams and the standard approach is to measure the frequency of these elements, or of a group of words of known sentiment orientation, in a document belonging to a given class. Text based features can also be derived from an ad hoc lexicon built in a preliminary phase of the study, by means of thesauri or semantic dictionaries such as

² Translated from Italian. S stands for ‘System’, U for ‘User’

WordNet³. To improve the accuracy of the classification, BoW are usually enriched with additional features which may be based on the proximity between the items to classify [13], on an ‘ad hoc’ taxonomy [15] or on the relationship of every word with the previous or the next one, as they appear in the parsing tree for the complete sentence [16].

When attempting to recognize opinions within a dialogue interaction rather than from analysis of a single text, more information about the context in which a sentence was pronounced is available. As we will see, on one hand this information makes recognition a more complex task but, on the other one, it provides more opportunities for a correct solution.

Let us consider the various forms of ‘context’ that occur in the interpretation of a dialogue move:

4.1.1 Local context: the user move

Most work on sentiment analysis was developed on monologs, such as movie reviews. Extending these methods to the analysis of single sentences or brief dialogue turns is not immediate. At a first glance, sentiment analysis should work well also in these cases: we might think to simply look at the prior polarity of subjective words such as ‘correct’ to interpret the polarity of sentences like ‘*I think my dietary habits are correct*’. However, after looking in our corpus we noticed that, to recognize the polarity of the user move, many other things have to be considered. In general, a word based approach demonstrated to be not powerful enough, especially in non binary classification tasks [16]. Mullen and Malouf [12], e.g., tried to identify the political affiliation of bloggers by analyzing their post on a web forum; purely text-based methods produced, in that case, a low accuracy because most posters from across the political spectrum used common terms such as ‘*gun control*’ or ‘*abortion*’, regardless of their opinion on those particular issues. These authors concluded that the accuracy could be improved by introducing rules based on the observation of how posters interact with each other, that is by adding information about the *context* in which a post is added to a discussion.

The situation becomes more complex if we want to perform sentiment analysis at the phrase level (short dialogue turns): the majority of problems is related to stop-words elimination involved by the BoW representation [11]. The prior polarity of words can be affected by linguistic factors that modify their ‘*contextual polarity*’ [16]. A typical example is the presence of negations, that may be local (‘*I think my dietary habits are not correct*’) or may involve longer-distance dependencies (‘*I don’t think my dietary habits are correct*’). If we simply rely on a word based approach, we might classify as identical the opposite cases ‘*I think my dietary habits are correct*’ and ‘*I think my dietary habits are not correct*’. This problem is due to the stop words elimination and has an impact on the recognition of the ‘strength’ of the opinion expressed: since adverbs are usually taken as stop words, sentences like ‘*I think my weight is pretty good*’ and ‘*I think my weight is really good*’ would be considered as identical.

In [9], the modifiers that change the semantic orientation (negations) of a term or its weight (intensifiers and diminishers) are named *valence shifters*. The cited paper presents a comparison between two approaches: in the first one, positive and negative terms in a document are counted, and the text is classified as having a positive orientation if more positive than negative terms are found (and vice versa) or neutral when the number of positive and negative terms is the same. Polarity of single terms is decided according to a dictionary. The second method takes into account contextual valence shifters in determining the semantic orientation of non-neutral words. A parser is used to determine which modifiers to apply to which terms. The term-counting method has the advantage of not requiring any training phase, since one can

simply rely on a lexicon established a priori: however, methods based on shifters evaluation proved to be more effective in terms of accuracy. The case of negation and, in general, of all modifiers, is also discussed in [15]: these authors present a new method for sentiment analysis based on extracting and analyzing *adjective appraisal groups* such as ‘*really good*’ or ‘*not so bad*’. Appraisal groups include an *head adjective* and an optional list of *appraisal modifiers* with nested scope, each denoting a transformation of one or more appraisal attributes of the head. Four attributes are used to describe every group: *attitude*, which gives the type of appraisal being expressed, *orientation*, which is the polarity (positive or negative) of the appraisal, *graduation*, which is the intensity of the appraisal and its focus and *polarity*, which says whether the group is *marked* as scoped in a polarity marker such as a negation. This taxonomy was employed to tag the lexicon in an enriched BoW representation in which terms were located in the four dimensional space by giving a value to all appraisal attributes.

There are also cases in which investigating the role of modifiers is still not enough. A typical example is: ‘*I can’t resist to a delicious sweet, what should I do?*’. In this example, lexicon with prior positive polarity prevails (‘*delicious sweet*’) and the action of modifiers (‘*I can’t resist*’) do not necessarily produce a negative classification of the turn: on the contrary, the negation of the verb strengthens the appeal of the ‘sweet’ word. In cases like this, the parsing tree of the sentence should be explored to capture its real semantics by analyzing the syntactic role of every word.

4.1.2 Wider context: dialogue pairs

In all the examples we saw so far, the context to consider in sentence interpretation was limited to a single user move. In other cases, however, knowledge of the previous system’s move is essential to recognize the user’s expression of opinion (see, e.g., the pair n.2 in Table 1). In our corpus, we found complete expressions of opinion like ‘*I think my weight is correct*’, but also several sentences such as ‘*pretty good*’ or ‘*I think it is ok*’ after the question: ‘*What do you think of your dietary habits?*’. In these cases, sentiment analysis may classify the user answer as generically positive or negative, but only thanks to our knowledge about the context we may say something about the user opinion. Beliefs inferred in the two cases have not the same level of validity: we will name *direct beliefs* those inferred from direct declarations of opinions, and *from-answers beliefs* those inferred from answers to system questions. Although they represent alternative ways of expressing beliefs, the first one is likely to provide a stronger evidence than the second one. An example: if (as in pair n.1) the system’s question was “*Do you think your diet is correct or would you like to change your dietary habits?*” and the user answers “*I think I should control my consumption of sweets*”, we may infer the user’s negative opinion about his own behavior with a lower level of certainty than if the question simply was: ‘*Tell me something about your diet*’. Strengthening or weakening of the level of certainty about an inferred belief may occur by combining different parts of a given move. For instance, in the pair n 3, the final user question ‘*what should I do?*’ strengthens the presumed U’s negative opinion about being overweight that was expressed in the first part of the sentence.

4.2 Progressively inferring beliefs: context is the whole dialogue

The problems discussed so far are only related to the task of determining the sentiment orientation of an individual user move and inferring a presumed belief from that local analysis. However, beliefs cannot be directly inferred from a unique expression of opinion. Recognizing the user beliefs relies on consideration of

³ <http://wordnet.princeton.edu/>

other aspects as well, such as the opinion holder, his status (how much credible, how much competent in the domain he is etc). One might express a personal opinion ('I think I'm drinking too much'), refer others' opinions ('My wife says I drink too much') or ask a question to the system playing the role of an expert in the domain, in order to check whether its beliefs are aligned with his own beliefs ('Do you think that drinking four beers a day is too much?'). In the first case, as we said, inference of beliefs from opinion expression is more direct and stronger, while in the second and the third one it is more indirect and weaker. Once we understand that in the sentence 'My wife says I drink too much' the opinion holder is U's wife, we need to know whether U thinks that his wife is credible and competent in the domain or whether he thinks she is (for instance) too anxious or oppressive: overall, the level of certainty of this belief depends on whether U considers that source as 'believable'.

In the third example ('Do you think that drinking four beers a day is too much?'), a question to the system may be interpreted in

terms of a condition of doubt rather than of a clear belief. We call *from-question* all beliefs generated by this kind of situation, *indirect* the kind of beliefs that originate from referring an external source's declaration rather than a personal opinion and *from-answer* the kind of beliefs that are inferred from answers to system questions. Figure 1 synthesizes the difference in inference of direct, indirect, from-question and from-answer beliefs, in context-based sentiment analysis. In this figure, 'z' represents a generic fact about the user diet; for instance: 'U is overweight', Overweight(U). The node '(Say U z)' represents a declaration of the type 'I am overweight'. The node '(Answer U z)' represents an answer 'No' to the system question 'Do you think your weight is correct?'. The node '(Say U (Say A z))' represents a declaration of the type 'My wife says my weight is not correct' and the node '(AskWhether U z)' represents a question like 'Do you believe that 90 kilos are too much for a person of my height?'




Figure 1. From (uncertain) opinion recognition to (uncertain) belief interpretation

As we said in Section 2, a belief may be inferred gradually from a cumulative expression of consistent opinions, and this inference process can be based on a mapping with the system's belief structure. A 'correct' behavior is the result of a number of components: in the case of healthy dieting, a good proportion of

vegetables, a right balance of the other components, regularity of meals and so on. Figure 2 represents the relationship about believing that own dietary behavior is balanced (or not) and believing that the components of dietary behavior are correct (or not).

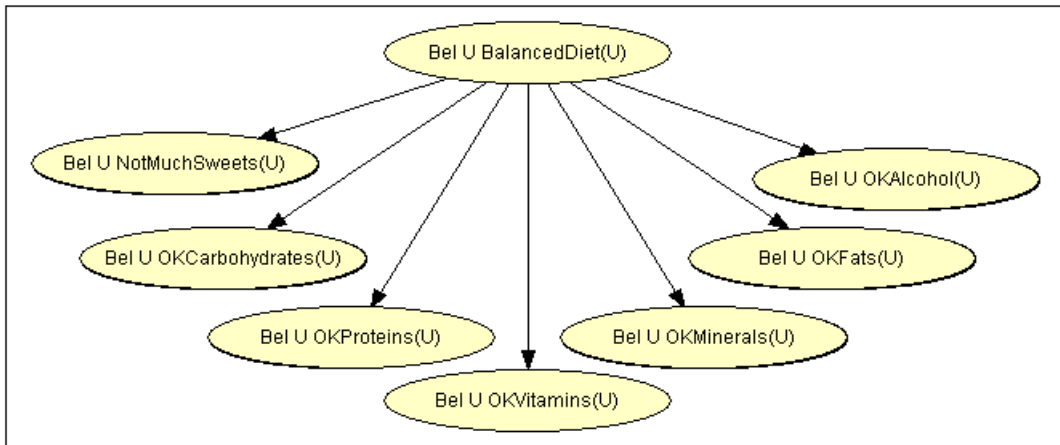


Figure 2. Relationship between generic and specific beliefs about own diet

A system playing the role of an advice-giver in this domain holds its knowledge in a 'consistent' set of beliefs. Recognizing how much consistent the user's set of beliefs appears to be is a dynamic process: the system progressively builds an image of the user mind by updating it after recognizing every expression of opinion, and by considering the strength and uncertainty of opinions expressed. Figure 3 represents the dynamic updating of the system's image of the user's beliefs during the dialogue. In this oriented graph, the relationship between every leaf node and its

parent node at time t ($\text{Bel U CorrectDiet}(U) t$) is a function of how important is the variable associated with the child node in defining a diet as 'correct'. The relationship between this last node and its parent nodes represents, in its turn, two effects: i) the progressive refinement of the system's image of the user's mental state, based on the information acquired during the dialogue and ii) the possible change in the user's belief about his own dietary behavior, produced by the system's suggestions and information provision.

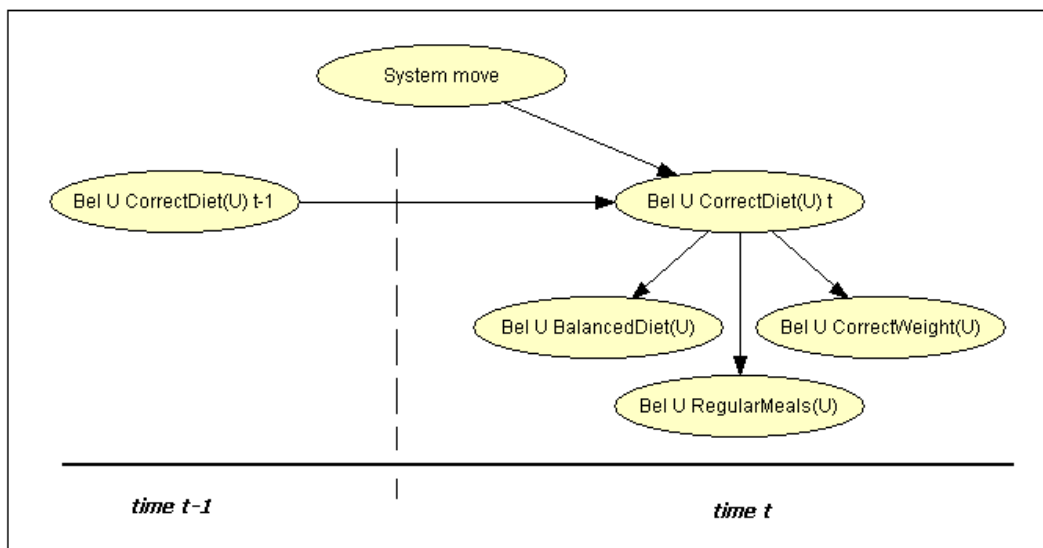


Figure 3. Dynamic updating of the user's set of beliefs

Table 2 describes how the system's image of the user beliefs evolves during our example dialogue, as soon as new information is acquired. Let us start from time t_1 (first dialogue pair). The sentence *'I think I should control my consumption of sweets'* is interpreted by the sentiment analyser as a direct statement of belief that he tends to take too much sweets in his diet (Say U MuchSweets(U)); this increases the probability of DirectBel U MuchSweets(U) (figure 1) and, consequently, decreases the likelihood that he believes his diet is well balanced ((Bel U BalancedDiet(U)): figure 2) and, therefore, correct (Bel U CorrectDiet(U)) (figure 3).

We now go to the next time slice in figure 3 (t_2). The sentiment analyser interprets the user move 'Yes' as a display of opinion that the present diet is not correct, although with a lower level of certainty than in the previous move (because it is a 'FromAnswer' type of belief). The probability of the corresponding node is updated... and so on. In this example, the system progressively acquires new information about the user during the dialogue, but apparently it does not influence the user mind with its moves, if not very slightly (as it just makes questions rather than giving overt suggestions).

Table 2. Progressive updating of the system's image of the user mind during the dialogue

N	Adjacent pair	Recognition	U's beliefs at time t_i
1	S: Do you think your diet is correct or would you like to change your dietary habits? U: I think I should control my consumption of sweets.	DirectBel U MuchSweets(U)	↓ Bel U CorrectDiet(U), t_1
2	S: Have you ever desired to change your diet? U: Yes	FromAnswerBel U not CorrectDiet(U)	↓ Bel U CorrectDiet(U), t_2
3	S: Do you think your weight is correct or would you like to change it? U: I am 4 kg overweight so I would like to change it. What should I do?	FromAnswerBel U not CorrectWeight(U)	↓ Bel U CorrectWeight(U), t_3
4	S: Why you never tried to do it? U: Because I'm not so patient	FromAnswerBel U not Enduring(U)	↓ Bel U ConditionsToChange(U), t_4
5	S: Do you think you assume all the substances you need? U: I regularly assume carbohydrates, proteins and vitamins. I probably assume not enough minerals and too much fat.	DirectBel U OKCarbohydrates(U) DirectBel U OKProteins(U) DirectBel U OKVitamins(U) DirectBel U not OKMinerals(U) DirectBel U not OKFats(U)	↑ Bel U BalancedDiet(U), t_5 ↑ Bel U BalancedDiet(U), t_5 ↑ Bel U BalancedDiet(U), t_5 ↓ Bel U BalancedDiet(U), t_5 ↓ Bel U BalancedDiet(U), t_5 overall: slight ↑ Bel U BalancedDiet(U), t_5
6	S: You don't have to completely eliminate fat foods. You have just to limit their proportions U: I can't organise my diet because I have no time.	DirectBel U not HasTime(U)	↓ Bel U ConditionsToChange(U), t_6
7	S: Sure, organize your diet could appear a challenging task especially when nobody can help you and you have to work or to study. U: I am 90% invalid	DirectBel U not IsAble(U)	↓ Bel U ConditionsToChange(U), t_7

Table 3 shows an excerpt of dialogue between a real human therapist and a subject with addictive behavior related to alcohol consumption [10]. The example shows how an example of successful persuasion strategy: the user U gradually moves from the 'precontemplation' stage (move 1), to the 'preparation' one (presumed, as it appears from move 7), passing through the 'contemplation' stage in the central part of the dialogue in which the awareness of adopting a wrong behavior gradually emerges, thanks to the ability of

the therapist in formulating 'ad hoc' questions. Let's denote with S a dialogue system equipped to emulate this behavior: in this example the variation in the system's image of the user belief changes gradually because the user is progressively persuaded by the system's suggestions; the model is identical to the one applied in the previous example but the node named 'SystemMove' in figure 3 contributes in this case to increase the probability of the node Bel U CorrectDiet(U), t .

Table 3. Progressive change of the user's belief during the dialogue, due to the persuasive strategy adopted by the advisor

N	Adjacent pair	Recognition
1	S: So one thing you've noticed is that you are drinking more now than you used to. What else? U: I can't really think of anything else. It doesn't really affect that much. I don't really get drunk very often.	U declares that he is not concerned about his drinking behavior.
2	S: So, although you know that your drinking has gone up over the past few years, it doesn't really seem to affect you more. U: Right. I can drink all night and it doesn't make me drunk. Other guys have trouble keeping up with me.	U declares he has not problems related to alcohol consumption
[...] U talks about his father's drinking and problems related to that behavior.		
3	S: Is there anything else you've noticed, any other way in which your drinking seems like your father's? U: Lately, there has been some times when I can't remember things that happened. I'll be drinking at a party, and the next morning I can't remember getting home. It's not too pleasant to wake up and have no idea where you left your car.	U acknowledges that he has got memory problems due to alcohol consumption.

N	Adjacent pair	Recognition
4	S: That can be scary, especially the first few times it happens. Give me an example. U: About 2 weeks ago, I was out with Bob and I guess I drank a little more than usual. When I woke up in the morning, I couldn't think of where my car was. I looked out the window and my car was in the driveway, and I guess I drove it there. I felt terrible.	
[...]		
5	S: Your situation doesn't seem bad to you. U: No, it doesn't. I've quit drinking for weeks at a time with non problem. And I have a couple of drinks and leave it alone. I have a good job and a family. How could I be an alcoholic? [...] I mean, I've got some problems, but I'm not a drunk.	U recognizes that he has got some problems with alcohol but declares that he <i>is not a drunk</i>
[...] <i>the therapist shows him results of blood tests and explains him how drinking might affect his health.</i>		
6	U: So I'm driving around legally drunk three times a week? So I have a higher risk, then?	U recognizes his problem behavior with alcohol
7	S: That's it U: I guess I have to do something about my drinking – either cut it down or give it up.	U declares he want to reduce or even stop alcohol consumption.

5. OUTLINE OF THE ALGORITHM

The analysis exposed so far suggested us how to define a markup language for isolating those dialogue turns in which opinions are expressed and from which to start for the definition of a method for automatically recognizing them. The WoZ corpus was annotated by three independent raters, by taking adjacent pairs as units of analysis (couples of adjacent System-User moves in the dialogue, as shown in our examples).

Every unit is labelled multiply with the following tags:

- *opinion polarity*: this tag may take a 'positive' or 'negative' value according to the polarity of the opinion expressed, or may be 'neutral' if no opinion is expressed;
- *opinion object*: this tag represents the aspect of the diet (or the examined behavior in general) which is the object of opinion expressed (consumption of sweets, carbohydrates, vegetables etc...);
- *move type*: this label specifies whether the opinion is expressed by means of a direct sentence ('direct opinion'), a 'question/answer to the system' ('from question', 'from answer' opinion) or by referring to a third person's opinion ('indirect opinion');
- *believability* of the opinion holder: this tag is used in case of 'indirect opinion' and may assume a value in the {'high', 'low', 'neutral'} set.

Results of the markup experiment showed that data are very sparse: therefore, we could not rely on classical machine learning techniques for automatically inferring tag values during the dialogue. Our idea is to combine sentiment analysis techniques, applied at the level of dialogue turns for opinion extraction, with decision rules, based on information about the context and the dialogue history, as observed in our corpus.

We sketched an algorithm which describes the dynamic recognition process we intend to implement: this algorithm is repeated every time a new user move is entered and includes the following steps:

1. A new user move is entered and treated as input for a module which implements sentiment analysis techniques. The process of opinion extraction exploits information about the context, as explained above, and gives as output the opinion polarity and its object (typically one of the aspect of the behavior considered, as shown in fig. 2);
2. The system infers a possible belief from the output produced at the previous step. The belief recognized

could be 'direct', 'indirect' fromQuestion' and 'fromAnswer'; each of them has a different weight in updating the set of user's beliefs. Updating of the particular belief inferred has effect on more general ones, as shown by the example in tab. 2;

3. At every dialogue move at time t , the system updates the image of the user mind on the basis of the new knowledge acquired at steps 1 and 2 and of knowledge at previous slice time $t-1$.

6. CONCLUSION

This contribution is a preliminary statement of the direction in which we are moving in our study about the relation between opinion expression and belief inference. The relationship between beliefs and action, attitudes and language is not so strict and in particular language is not easy to interpret. In this work, we have studied how beliefs can be dynamically inferred, during the interaction, from a set of consistent opinions in the scenario of a advice-giving dialogue system in a smart-environment.

Our long-term goal is to build a user-adapted dialogue system which dynamically fits persuasion plans to every specific situation and provides user-tailored suggestion about about healthy living habits.

According to Prochaska and Di Clemente's Transtheoretical Model of Change, we tried to define a method to automatically infer information about the beliefs related to the particular user's 'Stage of Change'. The main idea underlying our work is that the system may infer users' beliefs through a mapping with its own belief structure, by using as input for this process the expressions of users' opinions, as they can be observed in their linguistic behavior.

In this work we investigated the state of art in sentiment analysis techniques in order to find the main limitations that we have to cope with when operating in a dialogue context. The markup of our corpus showed us sparsity of data and this suggested us to sketch an algorithm which combines sentiment analysis for opinion extraction with decision rules based on the observation of the local or global context.

ACKNOWLEDGEMENTS

This work was financed, in part, by HUMAINE, the European Human-Machine Interaction Network on Emotion (EC Contract 507422).

References

- [1] J. Allwood (1981): *Language beliefs and concepts*. In Allwood, Fragsmyr, Hjort and Svedin (Eds): *Natural Resources in a Cultural Perspective*. FRN-Report 37, Stockholm.
- [2] P.R. Cohen and H.J. Levesque (1995). *Communicative actions for artificial agents*. Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95).
- [3] F. de Rosis, C. Castelfranchi, V. Carofiglio and R. Grassano (2003). *Can Computers deliberately deceive? A simulation tool and its application to Turing's Imitation Game*. *Computational Intelligence*, 19, 3.
- [4] F. de Rosis, N. Novielli, V. Carofiglio, A. Cavalluzzi and B. De Carolis (2006): *User modeling and adaptation in health promotion dialogs with an animated character*. *Journal of Biomedical Informatics, Special Issue on 'Dialog systems for health communications'*. T.Bickmore (Ed), 39, 5, 514-531.
- [5] C. Gauker: *Think out loud: An essay on the relation between thought and language*. Princeton University Press, 1994.
- [6] J. Gray and C. Brezeal (2005). *Toward helpful robot teammates: a simulation-theoretic approach for inferring mental state of others*. AAI Workshop on Modular Construction of Human-Like Intelligence, Pittsburgh.
- [7] P.H. Grice (1989). *Studies in the Way of Words*, Harvard University Press.
- [8] S. Gross (2006). *Can one sincerely say what one doesn't believe?* *Mind and Language*. 21, 1.
- [9] A. Kennedy and D. Inkpen (2006). *Sentiment classification of movie reviews using contextual valence shifters*. *Computational Intelligence*, 22, 2, 110-125.
- [10] W.R. Miller, S. Rollnick, *Motivational Interviewing – Preparing People to Change Addictive Behaviour*, The Guilford Press.
- [11] T.M. Mitchell (1997). *Machine Learning*. Boston, MA: McGraw Hill.
- [12] T. Mullen and R. Malouf (2006). *A preliminary investigation into sentiment analysis of informal political discourse*. AAI Symposium on Computational Approaches to Analyzing Weblogs, 159-162.
- [13] B. Pang and L. Lee (2004) *A Sentimental Education: sentiment analysis using subjectivity summarization based on minimum cuts*. Proceedings of the 42nd ACL, 271-278
- [14] J. Prochaska, C. Di Clemente and H. Norcross (2002) *In search of how people change: applications to addictive behavior*. *American Psychologist*, 47.
- [15] C. Whitelaw, S. Argamon and N. Garg (2005). *Using appraisal taxonomies for sentiment analysis*. Proceedings of the First Computational Systemic Functional Grammar Conference, Sydney.
- [16] T. Wilson, J. Wiebe and P. Hoffmann (2005) *Recognizing contextual polarity in phrase-level sentiment analysis*. Proceedings of HLT-EMNLP Vancouver