Enhancing Conversational Access to Information through a Socially Intelligent Agent

Berardina De Carolis, Irene Mazzotta, and Nicole Novielli

Abstract. Intelligent access to information could benefit of an effective and natural interaction metaphor. In this perspective, Embodied Conversational Agents (ECAs) can be seen as a promising approach to give to the user the illusion of cooperating with a partner rather than just using a tool. Embedding the HCI technology with human preferences and behavior justifies the attempt of implementing emotional and social intelligence aimed at exceeding the single ability to help the user. In this paper we present an ECA's architecture and methods useful to interpret the user attitude during her dialog with an ECA and behaving 'believably' in its turn. In particular, we present an agent architecture that is general enough to be applied in several application domains and that employs several ECA's bodies according to the context requirements.

Keywords: Natural Language Interaction, Conversational Access to Information, Emotional Intelligence, Embodied Conversational Agents.

1 Introduction

Intelligent Information Access has the main goal of providing a personalized access to information by exploiting information retrieval techniques. To this aim, the user behavior is observed, either directly or indirectly, in order to build user models which allow personalization (Kobsa 1993, Berkovsky et al. 2009).

We believe that a system providing intelligent and personalized information access, though, surely benefits of an intelligent presentation of the information content by exploiting methods for developing effective, usable and natural interaction metaphors.

An intelligent information system should therefore be equipped with and intelligent interface, that is an interface able to:

- adapt to the user;
- handle natural language dialogs using the appropriate strategies;
- decide, autonomously, when to activate themselves and how to respond to the (presumed) user needs;

Berardina De Carolis · Irene Mazzotta · Nicole Novielli
Dipartimento di Informatica, Università degli Studi di Bari,
Via Orabona, 4 - 70125 - Bari, Italy
e-mail: {decarolis,mazzotta,novielli}@di.uniba.it

G. Armano et al. (Eds.): Intelligent Information Access, SCI 301, pp. 1 – 20.
springerlink.com

- consider user emotions and social attitude during the dialogs and behave 'believably' in its turn.

In this perspective, Embodied Conversational Agents (ECA) can be seen as a new metaphor of human-computer ‘intelligent’ interaction which promises to be effective (Cassell and Bickmore 2003) if the hypothesis that ‘characters contribute to more sociable and user-friendly interfaces’ is taken for granted (Lee and Nass 1999). A well designed ECA should give the users the illusion of cooperating with a human partner rather than just ‘using a tool’. The more the agent succeeds in this goal the more the users are expected to attach some anthropomorphic features to them and to show signs of affective (emotional, social) involvement in the interaction. Therefore, in developing a "computer conversationalist" that is embedded in a ECA and that is able to exhibit these capabilities it is important to conceive its architecture so as to:

1. start from the interpretation of the spoken or written utterance;
2. reason on the various information the user intends to convey (emotion, social attitude, performative, content, etc.) and then to trigger communicative goals according to the current belief representation of the state of the world;
3. achieve these goals through a set of communicative plans (“what to say”) that can then be rendered as a combination of voice and animations of the agent’s body (“how to say”).

To achieve believable natural language conversations, ECA systems are quite complex and require the combination of several assemblies (Huang et al., 2008). In recent years, the members of the international research community have been jointly working towards the definition of a standard framework for the generation of behavior of virtual agents. The SAIBA project (Vilhjálmsson et al., 2007) represents one of the major efforts towards the unification of ECA standards and aims at unifying the key interfaces in the multimodal behavior generation process. In particular, the SAIBA architecture is structured on three different layers, each of them implemented to serve to a different function of the agent behavior planning and realization. Moreover, as far as affective computing is concerned, the SEMAINE API provides an open source framework for building systems embedding emotional intelligence: the SEMAINE system is a full-scale system resulting from the integration of several existing and new components (Schröder, 2010).

The example provided by these projects demonstrates how scalability and openness represent two extremely important challenges in developing architectures for believable ECAs: as far as scalability and openness are ensured, standard architectures can be easily extended to meet the specific needs of the various interaction scenarios and application domains (Bevacqua et al., 2009). The architecture proposed in the present paper perfectly fits in the vein of this ongoing research.

In particular, this paper describes our experience in the design and implementation of a scalable architecture of a believable ECA that interacts with the user for providing advices in a domain where considering social and affective factors is crucial. To this aim, the architecture is designed so as to dynamically model and build different agent’s functionalities, according to the application domain needs.

Moreover, taking into account context factors, it allows: i) employing the most appropriate dialog strategy (for instance, information giving vs. persuasion dialogs), ii) simulating more or less rational or affective agent's behaviors (Carofiglio et al. 2009), iii) using different agent's bodies.

The paper is structured as follows: Section 2 describes the agent architecture; then in Section 3 we provide an overview of the Interpersonal Stances Modeling; Section 4 presents a Model of Emotion Activation. In Section 5 we describe the Dialog Modeling component with an example of implemented dialog (Section 6). Section 7 shows a possible extension of our agent architecture, by describing how a persuasion module can be integrated and used opportunely in the agent architecture. Conclusions and future work directions are presented in the last Section.

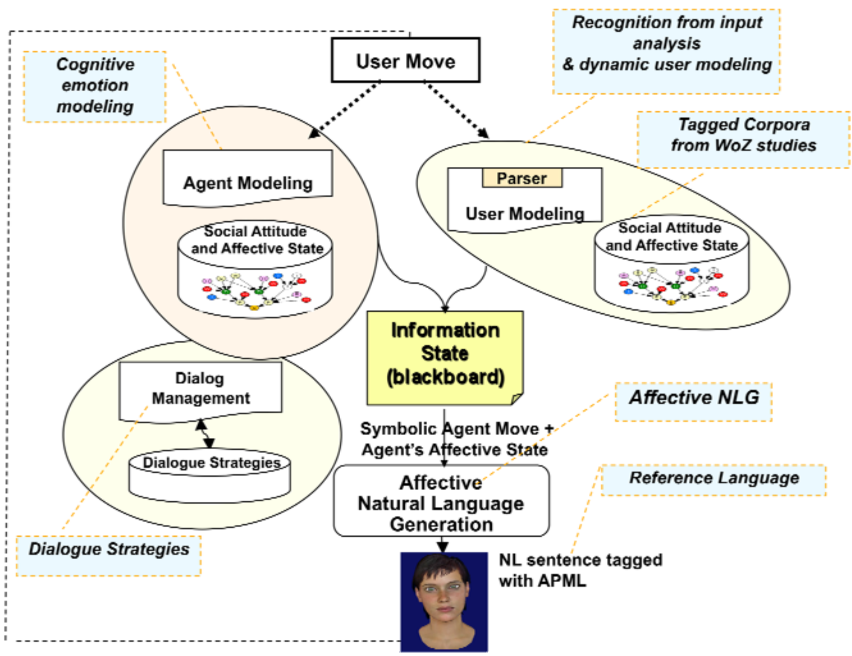

Fig. 1 An overview of the ECA architecture

2 The Architecture

The ability to exhibit an emotional state and/or social signs is a shallow form of the intelligence an agent can show. The recognition of the social attitude and of the emotional state of the interlocutor should be utilized to drive reasoning behind the dialog between the user and the ECA. This implies studying how these factors may affect the ECA architecture. In our opinion, when developing an ECA, the following issues should be addressed:

- the user move should be interpreted so as to detect, beside the linguistic content: i) which is the social attitude of the user and ii) which emotions arise during the dialog;
- how these factors influence the dialog course by changing the priority of communicative goals, dialog plan and surface realization of communicative acts.

Figure 1 illustrates the architecture we propose to handle these issues. This architecture has been conceived as composed by two main functional modules: the "mind" and the "body" of the agent.

2.1 The Mind

The user move is a rich information source that allows extracting knowledge about the user's intention, her social attitude, emotional state, and so on. In our approach, the "mind" of the agent uses two main different knowledge sources for reasoning on the user move and then formalizing beliefs that are useful for planning its dialog move: the user and the agent models.

The user model component allows to reason on the user's beliefs (i.e. the user move "I love fruit!" will be transformed into the correspondent belief that can be used to adapt the dialog strategy) and on the user's social attitude during the dialog (i.e. the user move "It's nice to talk to you!" will be interpreted as a sign of friendly disclosure towards the ECA). While beliefs on knowledge, preferences and interests of the user are inferred according to an approach previously employed in another system (de Rosis et al. 1992), in this paper we will explain how the user social attitude is recognized and monitored with a dynamic model based on Belief Network (DBN) (Jensen 2001).

The agent model is also based on a DBN which mainly aims at triggering emotions that arise in the agent mind during the interaction, in a given situation, according to the agent's personality and to the social context in which the dialog occurs. Starting from what has been inferred by the user model component and from the emotions triggered in the mind of the agent, the dialog management module computes the agent move using a strategy that will be explained later in the paper.

The information exchange among these modules is managed using a common blackboard called information state (Larsson and Traum 2000). It represents the memory of the agent and stores beliefs about the current state of the dialog, the dialog history, the current dialog move and the move scheduled for execution. This approach allows employing different methods and techniques giving to the architecture a degree of openness and scalability.

2.2 The Body

While the move computed by the "mind" module contains the meaning to express ("what to say"), the "body" has to convey these meaning according to its communicative capabilities ("how to say"). In order to decouple meanings from signals we use a mark-up language: APML (De Carolis et al. 2004). These meanings include the communicative functions that are typically used in human-human dialogs: for instance, syntactic, dialogic, meta-cognitive, performative, affective, deictic, adjectival and belief relation functions (Poggi et al. 2000).

The use of a reference language gives the possibility to employ different bodies and different platforms and devices without changing the mind of the agent. In fact, in order to express the same meaning using different signals according for instance to the context or to the capabilities of the body of the employed ECA, each ECA's body has a conditional meaning-signal table that allows to appropriately translate an APML tag into tags expressed in Signal Expression Markup Language (SEML). SEML tags define the expressions that can be performed on each channel of the Body as described in (De Carolis 2005).

Let's see now in more details how these modules work.

3 Interpersonal Stances Modeling

After several forms of 'anthropomorphic behavior' of users towards technologies were demonstrated (Reeves and Nass 1996), various terms and concepts have been employed to denote this behavior and describe it. Paiva (2004) talks about empathy, Hoorn and Konijn (2003) address the concept of engagement, involvement, sympathy and their contrary, distance. Cassell and Bickmore (2003) adopt the Svennevig's theory of interpersonal relations.

We refer to Scherer's concept of interpersonal stance as a category which is "characteristic of an affective style that spontaneously develops or is strategically employed in the interaction with a person or a group of persons, coloring the interpersonal exchange in this situation (e.g. being polite, distant cold, warm, supportive, contemptuous)".

In particular, in referring to the social response of users to ECAs, we distinguish warm from cold social attitude, according to the Andersen and Guerrero's definition of interpersonal warmth (Andersen and Guerrero 1998) as "the pleasant, contented, intimate feeling that occurs during positive interactions with friends, family, colleagues and romantic partners".

We studied this attitude and the factors affecting it (Novielli et al., 2010) by observing the verbal and prosodic behavior of 60 subjects interacting with an ECA in a Wizard of Oz simulation study (de Rosis et al. 2007). More details about the WoZ study may be found in (Clarizio et al. 2006). In particular, we defined a markup language (Table 1) for the user moves after carefully examining our corpus and considering suggestions from the studies about verbal expression of social attitude (Andersen and Guerrero 1998, Polhemus et al. 2001, Swan 2002). Dynamic recognition of these individual signs during the dialogue enables not only to estimate the overall social attitude value but it also allows the agent to adapt its dialogue plan accordingly: for example, if the user tends to talk about herself, in the following moves the ECA will use this information to provide more appropriate suggestions. The overall social attitude of the user will be inferred dynamically from the history of the signs recognized during the dialogue to adapt the ECA's language style, voice and facial expression.

Table 1 Linguistic signs of Social Attitude and their definition

Linguistic Signs of Social Attitude with definition
Friendly self-introduction: The subjects introduce themselves with a friendly attitude (e.g. by giving their name or by explaining the reasons why they are participating in the dialogue).
Colloquial style: The subject employs an informal language, dialect, proverbs
Talks about self: The subjects provide more personal information about themselves than requested by the agent.
Personal questions to the agent: The subject tries to know something about the agent's preferences, lifestyle etc., or to give it suggestions in the domain.
Humor and irony: The subjects make some kind of verbal joke in their move.
Positive or negative comments: The subjects comment the agent's behavior, experience, domain knowledge, etc.
Friendly farewell: This may consist in using a friendly farewell form or in asking to carry-on the dialogue.

Three PhD students labeled independently the corpus of WoZ dialogues with our markup language. According to the result of the annotation experiment we defined a set of linguistic cues that could be considered as salient (Lee et al. 2002) for every given of social attitude. These cues are organized into semantic categories. Every new user move is categorized as 'showing a particular sign of social attitude' if it includes some word sequences belonging to semantic categories which are defined as 'salient' for the considered sign (Novielli et al., 2010). Recognition of linguistic signs of social attitude is performed by using Bayesian classification and can be enriched with acoustic analysis of user move, as described in (de Rosis et al. 2007).

3.1 Dynamic Modeling of the User Attitude

The user modeling procedure integrates (i) language analysis for linguistic cues extraction and (ii) a dynamic belief network (DBN) which considers the context in which the move was uttered. DBNs (Jensen 2001), also called time-stamped models, are local belief networks (called time slices) expanded over time; time slices are connected through temporal links to constitute a full model. The method allows us to deal with uncertainty in the relationships among the variables involved in the social attitude estimation (Table 2). The DBN formalism is particularly suitable for representing situations which gradually evolve from a dialog step to the next one. We applied results of the corpus analysis to learn from the annotated data a model of the user's mental state (Carofiglio et al. 2005) which includes the dimensions of interest for dialog adaptation. In particular: in learning the temporal part of our DBNs, we took every single user move in the corpus as an independent observation and applied the K2 algorithm (Cooper and Herskovitz 1992); in learning the temporal link between the monitored variable *Satt* at two subsequent time instants, we took every dialog as an observation to measure the conditional probability that *Satt* takes a given value at time t, given its value at time t-1.

The DBN (Figure 2) is employed to infer how the social attitude of the user evolves during the dialog in relation to the dialog history. The social attitude is the hidden variable of our model, that is the variable we want to monitor, which depends on observable ones, such as the 'stable' characteristics of the users (their background and gender), the context in which the move was entered (previous agent move) and the linguistic features of the user move recognized by our Bayesian classifier (leaf nodes of our DBN). Intermediate variables represent the signs of social attitude listed in Table 1.

Links among variables describe the causal relationships among stable characteristics of the users and their behavior, via intermediate nodes. DBNs, as employed in this paper, are said to be 'strictly repetitive models'. This means that structure and parameters of individual time slices is identical and temporal links between two consecutive time slices are always the same. We use a special kind of strictly repetitive model in which the Markov property holds: the past has no impact on the future given the present. In our simulations, every time slice corresponds to a user move, the stable user characteristics do not change from time to time (this is why we omitted the nodes Back and Gend from the figure) and temporal links are established only between dynamic subject characteristics in two consecutive time slices.

Table 2 Variables of our model

Variable category	**Variable Name**	**Label**
Stable user characteristics	Background	Back
	Gender	Gend
Context	Last agent move type	Ctext
	User move type	Mtype
Monitored variable	User attitude towards the agent	Satt
Signs of social attitude	Familiar style	Fstyl
	Friendly self-introduction	Fsint
	Talks about self	Perin
	Question about the agent	Qagt
	Friendly farewell	F-Fw
	Comments (positive and negative)	Comm
Result of linguistic analysis	Cues of familiar style	Pfstyl
	Cues of friendly self-introduction	Pfsint
	Cues of talks about self	Pperin
	Cues of questions to the agent	Pqagt
	Cues of friendly farewell	Pffw
	Cues of comments	Pcomm

At the beginning of interaction, the model is initialized by assigning a value to the stable user characteristics (e.g. female user with background in Humanities). At every dialog step, knowledge about the context and evidence produced by linguistic analysis are entered and propagated in the network: the model revises the probabilities of the social attitude node. The new probabilities of the signs of social attitude and stage of change are used in formulating the next agent move, while the probability of the social attitude node supports revising high-level planning of the agent behavior.

We performed an evaluation of the model to examine how a variation in the threshold of the probability of the monitored variable (Satt) affects sensitivity and specificity of the model in recognizing this feature. For more details about model validation please refer to the study described in (Clarizio et al. 2006).

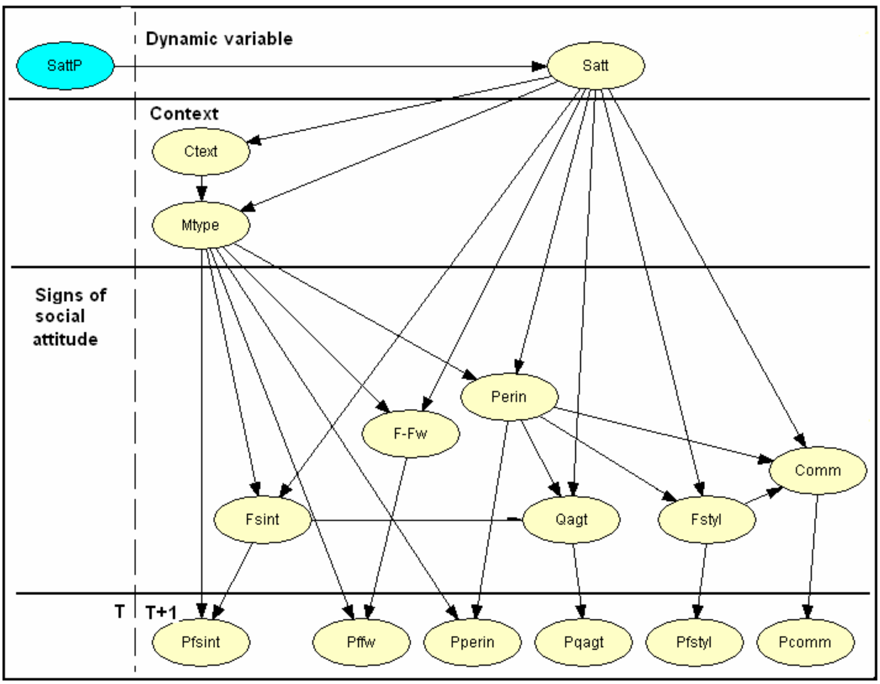

Fig. 2 User Model for the Social Attitude, a generic time-slice

4 A Model of Emotion Activation

In our emotion modeling method (de Rosis et al. 2003) we pay particular attention to how emotions change of intensity with time, how they mix up and how each of them prevails, in a given situation, according to the agent's personality and to the social context in which the dialog occurs. So far, we focused our attention on event-driven emotions in Ortony, Clore and Collin's (*OCC*) theory (Ortony et al. 1988). In this theory, *positive* emotions (happy-for, hope, joy, etc.) are activated by *desirable* events while *negative* emotions (sorry-for, fear, distress, etc.) arise after *undesirable* events.

Events concerning the agent are in the *Well-being* category (joy, distress), events concerning other people are in the *FortuneOfOthers* category (happy-for, sorry-for, envy and gloating) while future events are in the *Prospective* category (fear, hope). In Oatley and Johnson-Laird's theory, positive and negative emotions are activated (respectively) by the belief that some goal will be achieved or will be threatened (Oatley and Johnson-Laird 1987). A cognitive model of emotions that is built on this theory should represent the system of beliefs and goal behind emotion activation and endows the agent with the ability to *guess the reason why she feels a particular emotion and to justify it*. It includes the ingredients that enable representing *how the Agent's system of goals is revised* when emotions are felt and how this revision influences planning of subsequent dialog moves.

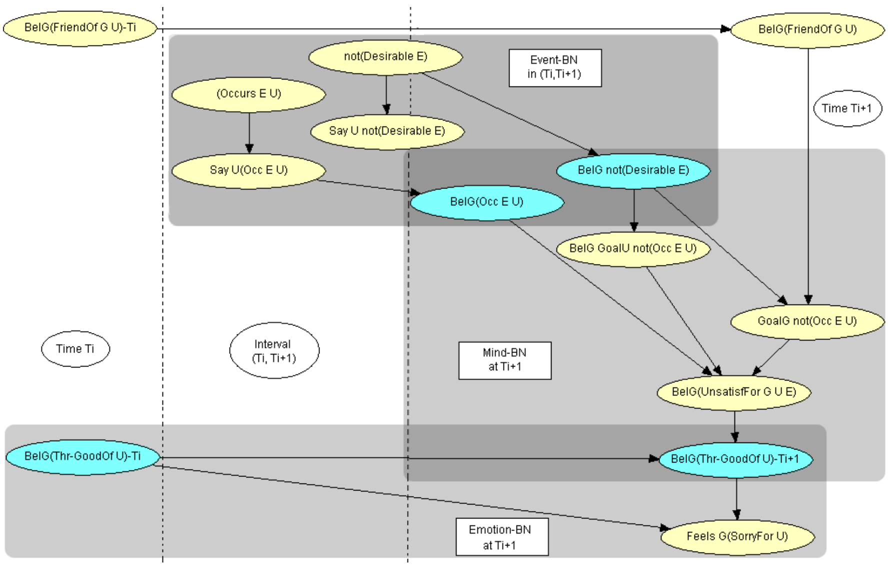

Fig. 3 A portion of the DBN that represents the agent's mental state showing the triggering of Sorry-For

Our model of emotion activation is represented with a *DBN* (Jensen 2001). We use *DBN*s as a goal monitoring system that employs the observation data in the time interval (T_i, T_{i+1}) to generate a probabilistic model of the agent's mind at time T_{i+1}, from the model that was built at time T_i. We employ this model to reason about the consequences of the observed event on the monitored goals. We calculate the intensity of emotions as a function of the *uncertainty* of the agent's beliefs that its goals will be achieved (or threatened) and of the *utility* assigned to achieving these goals. According to the utility theory, the two variables are combined to measure *the variation in the intensity of an emotion* as a product of the change in the probability to achieve a given goal, times the utility that achieving this goal takes to the agent (Carofiglio et al. 2008).

Let us consider, for instance, the triggering of *sorry-for* that is represented in Figure 3. This is a negative emotion and the goal that is involved, in this case, is *preserving others from bad*. The agent's belief about the probability that this goal will be threatened (Bel G (Thr-GoodOf U)) is influenced by her belief that some undesirable event E occurred to the user (BelG(Occ E U)). According to Elliott and Siegle (1993), the main variables influencing this probability are the desirability of the event (Bel G not(Desirable E)) and the probability that the agent attaches to the occurrence of this event (Bel G (Occ E U)). Other factors, such as the social context (Bel G FriendOf G U)), affect the emotion intensity. The model of the agent state at time T_{i+1} is built by automatically combining several *BN*s: the main one (Mind-BN) and one or more Event-BNs and Emotion-BNs. In the Event-BNs, the user moves are interpreted as *observable* consequences of occurred events that activate emotions through a model of the impact of this event on the agent's beliefs and goals. The strength of the link between what the user said (Say U (Occ E U)) and the *hidden* event (Occ E U) is a function of the user sincerity; the link between this observation and the agent's belief (Bel G (Occ E U)) is a function of how believable the agent considers the user to be. Therefore, the more sincere the user is and the more likely the event is a priori, the higher will be the probability that G believes in the occurrence of the event E. Similar considerations hold for the evaluation of how *desirable* the event is (Bel G (Desirable E)); these nodes are leaves of the Event-BN. They are, as well, roots of Mind-BN: they influence G's belief that U would not desire the event E to occur (Bel G Goal U ¬(Occ E U)) and (if G is in an *empathy* relationship with U and therefore adopts U's goals), its own desire that E does not occur (Goal G ¬(Occ E)). This way, they concur to increase the probability that the agent's goal of *preserving others from bad* will be threatened.

Variation in the probability of this goal activates the emotion of *sorry-for* in G through the Emotion-BN. The intensity of this emotion is the product of this variation times the *weight* the agent gives to the mentioned goal. According to Carbonell, we define a personality as a cognitively plausible combination of weights the agent gives to the goals represented in the model (Carbonell 1980).

The strength of the link between the goal-achievement (or threatening) nodes at two contiguous time instants defines the way the emotion, associated with that goal, decays, in absence of any event influencing it. By varying appropriately this strength, we simulate a more or less fast decay of emotion intensity. Different decays are attached to different emotion categories (positive vs. negative, FortuneOfOthers vs. Wellbeing and so on) and different temperaments are simulated, in which the *persistence* of emotions varies. The agents' affective state usually includes multiple emotions. Different emotions may coexist because an event produced several of them at the same time or because a new emotion is triggered while the previous ones did not yet decay completely. We describe in (Carofiglio et al. 2008) how we modeled the two mentioned mixing metaphors (*microwave oven* and *tub of water*, in Picard's terminology (Picard 1997)).

5 Dialog Modeling

The dialog manager includes three main layers:

1. a *Deliberative layer* that selects the goal with the highest priority and the correspondent plan and stores in the agenda the actions of the plan;
2. a *Communicative layer* that executes the next action in the agenda;
3. a *Reactive layer* that decides whether the goal priority should be revised, by applying reaction rules.

The dialog manager, and in particular the deliberative module, decides what goals to trigger and to pursue during the dialog, starting from the interpretation of the user move in terms of content and social attitude and according to the emotion triggered in the agent mind.

As the dialog evolves these factors may change what has been planned at two levels:

- by manipulating the inner aspects of the emotional response of our agent with an algorithm of activation/deactivation of its goals and of dynamic revision of their priorities;
- by deciding whether the agent should manifest its emotion and how.

Handling these issues is the main task of the Reactive Layer. The idea is that the agent has an initial list of goals, each with its priority, some of which are inactive: every goal is linked, by an application condition, to a plan that the agent can perform to achieve it. The communicative actions correspondent to active plans are put in the agenda maintained by the information state. The agent starts the dialog by executing these actions but, as we said in the Introduction, the agent applies some form reasoning on the user move. The recognized social attitude and the emotion triggered in the agent mind are used to implement social and emotion-based dynamic revision of goals and consequently of the dialog.

To achieve this aim, the following knowledge sources are employed by the dialog management modules:

1. *Agent's beliefs* that regard:
   – *long-term settings* that are stable during the dialog and influence the initial priority of the agent goals and therefore its initial plan, initiative handling and behavior: agent's personality, its role, its relationship with the user;
   – *short-term settings* that evolve during the dialog and influence goal priority change and plan revision: in particular, the emotional state of the agent and the social attitude of the user.

2. *Agent's goals* can be in one of the following relations among themselves:
   – *Priority:* $g_i < g_j$: g_i is more important, to the agent, than g_j. If this relation holds and no constraints or preconditions are violated by satisfying it, g_i will be achieved before g_j.

- *Hierarchy:* $H(g_i, (g_{i1}, g_{i2}, \ldots, g_{in}))$: the complex goal *gi* may be decomposed into simpler subgoals $g_{i1}, g_{i2}, \ldots, g_{in}$, which contribute to achieving it.
- *Causal Relation:* $Cause(g_i, g_j)$: the plan achieving the source goal g_i is a precondition for the plan achieving the destination goal g_j.

3. *Plans* that are represented as context-adapted recipes; a recipe may be applied when some preconditions hold; its application affects the dialog state (agent's and user's mental state and interaction settings). In the healthy eating domain, our agent adopts the typical plan of intelligent advice systems:
   - *situation-assessment*, to acquire information about the user,
   - *describe-eating-problems*, to describe eating problems and their possible origin,
   - *suggest-solution*, to describe how to eat better and to overcome problems,
   - *persuade-to change*, to convince the users to change their eating habits.

4. *Reaction rules* that implement goal-revision strategies. They may produce, in general, the following effects on the dynamics of plan activation:
   - *add details* when the user asks for more information;
   - *reduce details* in case of urgency;
   - *abandon temporarily a plan* to activate a new subplan to reassure, motivate or provide more details;
   - *abandon a subplan* when its goal has been achieved: for example, when the user seems to know the information the agent is providing;
   - *substitute a generic subplan* with a more specific and situation-adapted one;
   - *revise the sequencing of plans,* to respond to the User request of taking the initiative. This is the most delicate situation: to be cooperative, the agent should leave aside its dialog plan and follow the user request; however, as we said, communicative goals may be linked by causal relations. Therefore, when the users show the intention to take the initiative in the dialog, the agent checks whether their goal may be activated immediately or whether some preconditions have first to be satisfied. It then satisfies these preconditions with the shortest subplan before satisfying the user request (De Carolis 1999).

As far as emotions and social factors are taken into account, according to Oatley and Johnson- Laird (1987) that claimed that *human plans are much more flexible than those so far explored in AI* , our reactive planning method takes these factors into account from two points of view:

1. *rules* regulating the *goal priority revision* by formalizing the following strategies:
   i. in case of *urgent events*, reduce the detail of information provided by upgrading the priority of ”most relevant” subgoals and downgrading the priority of details;
   ii. in case of *desirable or undesirable events* occurred to the user, display *altruistic* social emotions (sorry-for and happy-for) by means of ”full

expression" goals, that is by verbal and nonverbal means, and give them the highest priority; revise the priority of other goals; hide *egoistic* social emotions as envy and gloating;

iii. in case of *desirable events* occurred to the agent, activate surface expression goals: use verbal and nonverbal means to express them but leave the priority of other goals unvaried;

iv. in case of undesirable events (again occurred to the agent), activate behavior control goals: avoid displaying any emotional reaction by activating, at the same time, repair goals.

With these rules, we formalize a situation of empathic reaction in which the agent temporarily substitutes the presumed goals of the user for its own, when these goals are due to an emotional state of the user (Picard 1997). If an undesirable event occurs to the users, what they are presumed to need is to be convinced that the agent understands the situation and does its best to solve the problem. If something desirable occurs to them, they need to know that the agent shares their positive experience. If, on the contrary, the undesirable event does not concern the users, they probably want to be sure that this will not interfere negatively with the dialog.

2. *rules* regulating *emotion display*. Once felt, an emotion *E* may be hidden or displayed by the Agent *Ag*. Again, this decision (although not always taken consciously) may be influenced by personality factors and by the interaction context. In (De Carolis et al. 2001) we describe how this aspect of the emotional behavior of our agent can be modeled by means of rules that regulate activation of display goals.

   For example, the following rule:

```
If (Feel Ag E) ∧ Is-a (E WellBeing) ∧ (Valence E Nega-
tive) ∧ (Adoptive Ag U) then Goal G (¬ Display Ag E)
```

activates the goal *G* of hiding a socially negative emotion felt, because the agent has an adoptive relation with the user (i.e. fear, anger, etc.).

On the contrary, the following one:

```
If (Feel Ag E) ∧ Is-a (E WellBeing) ∧ (Valence E Posi-
tive) (Adoptive Ag U) then Goal G (Display Ag E)
```

activates the goal *G* of showing a positive emotion.

6 An Example of Dialog

To test how incorporating affective and social factors may contribute to the believability of the dialogue, we built a domain-independent simulator that uses the architecture described so far. The following (Figure 4) is an example dialogue in the application domain of healthy eating (Berry et al. 2005). In this dialogue the agent and user moves are denoted respectively with Vi and Uj. Vi moves are generated by our dialogue simulator, while Uj moves are introduced in free text.

V0: Hi, my name is Valentina! I'm here to suggest you how to eat well. What's your name?
U0: My name is Jana and I'm happy to interact with you. ***Friendly self-introduction***
V1: Good to hear your interest Jana! Hope you'll find our dialogue useful. How can I help you? ***ECA reciprocates the Colloquial attitude***
U1: I've been eating too much and badly in the past year and I don't feel comfortable with myself! ***Talk about self & Colloquial style***
V2: Oh! That's a bad situation indeed, but I understand! Let me try to see how can help you. Which is your age? ***ECA reciprocates the Warmth and shows a Sorry-for***
U2: I'm 17.
V3: Well: teens are notorious for being hungry frequently, because they need to eat more to support the major growth of their bodies. Maybe you only need an appropriate diet and change your eating habits.
U3: Yes, you are right! I definitely need a more regular regimen. ***Talk about self & Colloquial style***
V4: Good, let's try to do something then. Did you ever think to change your diet? ***ECA reciprocates the Warmth and shows a Happy-for***
U4: Yes, sometimes I did it and got very good results. Now, either because of my laziness or my stressing rhythms I can't lay myself into more regular habits. But now I want to try again! ***Talk about self & Colloquial style***
V5: Good, that's the way I want to see you! You seem to be in a more positive attitude now! ***ECA reciprocates the Friendly attitude & shows a Happy-for***

Fig. 4 An example of dialogue in the healthy eating domain

In the present prototype, translation of the user moves into a symbolic communicative act is quite primitive (just based on keyword analysis and on the dialogue context). For instance the U1 move 'I've been eating too much and badly in the past year and I don't feel comfortable with myself!' is interpreted as Say U (Occ EatingTooMuch U), Say U not(Desirable EatingTooMuch U). Symbolic communicative acts are inputs of the cognitive emotion model which, in this example, activates the Sorry-For. At the same time, linguistic cues of Colloquial Style and Talks about Self are detected and evidences about these signs contribute to increase the overall likelihood of observing a warm social attitude of the user. Hence, in the subsequent move (V2) the agent reacts by expressing her Sorry-For ('That's a bad situation indeed!') and by reciprocating the warm social attitude through the use of some small talk ('But I understand!').

The next move U2 does not show any particular sign of social attitude and does not provide any evidence which could potentially cause emotion triggering. Here the sorry-for decays due to the absence of any more stimuli. The dialogue goes on quite neutrally until the user claims her intention to change her diet, in U3. This event causes the triggering of a light Happy-For, whose intensity depends on the belief of the agent about the user sincerity, that is how true the agent beliefs the user wants to change her diet given that the user claimed it.

Then, the user reacts to the agent question by friendly talking about self. As a consequence, a higher level of the user social attitude is estimated, causing the agent to reply with a colloquial style in her next move ('Good, that's the way I want to see you!'). Moreover, the user states again her intention to change her diet causing an increase of the intensity of the Happy-For felt by the agent.

7 Exploiting the Potential of New Dialog Strategies

Providing an intelligent and conversational access to information, in our opinion, requires not only to dynamically adapt the information provision during the dialog but also to employ the most appropriate dialog strategy. An intelligent system should be able, for instance, to increase the user intention to accept the system suggestion/recommendation (Mahmood et al. 2009).

For instance, in the example dialogue (Figure 4), the persuasion attempt performed by the system is represented by a single dialogue move (V3), implementing the sub-plan 'persuade to change' in the scope of the overall dialogue plan of 'intelligent advice system'.

Though, users may rise objections or show perplexity during the dialog. Therefore the ECA has to reason in order to answer to the user reaction. For example, the following dialog excerpt (Figure 5) represents a variation to the one in Figure 4 in which the user objects to the ECA suggestion (U4).

The persuasive attempt in the dialogue in Figure 5 is generated by PORTIA (Mazzotta et al. 2007) a reasoning module able to decide on the most promising persuasion strategy to apply in a given scenario. It is a user-adapted persuasion module capable of simulating the persuasion process used by humans to convince someone to perform a given action. It mainly focuses on two typical aspects of the human persuasion in order to produce effective persuasion attempts in different contexts: on one hand, the ability of reasoning on the potential strength of alternative persuasive strategies for a given user, in order to select the most appropriate one; on the other hand, the capability of combining rational and emotional modes of persuasion. The system is based on the theory of *a-rational* persuasion (Miceli et al. 2006), and the strategies represented in the model are the result of a combination of theoretical (Walton 1992, 1996, Petty and Cacioppo 1896) and empirical (Mazzotta and de Rosis 2006) background. The key points of the system are the separation between *reasoning* and *argumentation* phases in the persuasion process (Walton 1990) and the use of *Belief Networks* to represent the uncertainty inherent in this form of practical reasoning (Pearl 1988).

V3: Well: teens are notorious for being hungry frequently, because they need to eat more to support the major growth of their bodies. Maybe you only need an appropriate diet and change your eating habits.
U3: How can I do?
V4: Why don't you try to eat more fruit and vegetables?
U4: I don't like them very much…
V5: Maybe you don't know that they have a lot of benefits on your health.
U5: I'm young and have a lot to do instead being worried for my health!
V6: Fruit and vegetables contribute to improve your appearance. FDA also says that they help you to have healthy skin and hair and it is an authoritative voice!
U6: But cooking vegetables is boring and I prefer spending my time among the people and making new friends.
V7: A dinner with fresh and tasty salads is easy to prepare and superb to spend good time with friends. I'm sure you can do it if you wish.
U7: Yes, you are right! I definitely need a more regular regimen.

Fig. 5 An example of the ECA persuasion attempt

The PORTIA's persuasion strategies (Mazzotta et al. 2008) are summarized in Table 3.

Table 3 A summary of the Persuasion Strategies used by PORTIA

<table>
<tr><td colspan="6">General induction of intentions strategy</td></tr>
<tr><td colspan="6">It may be summarized as follow: "If User has the goal g and he believes that doing the action a implies achieving g in a more or less near future, and he believes that has the ability to do a, then probably user intends to do a" (from Miceli et al. 2006).</td></tr>
<tr><td colspan="3"><u>Rational induction of intention</u>
It focuses on rational goals like 'to be in good health', 'to have a good appearance', …</td><td colspan="3"><u>Emotional induction of intention</u>
It focuses on rational goals like 'to make friends', to be in good mood', …</td></tr>
<tr><td colspan="6">Activation of goal strategy</td></tr>
<tr><td colspan="6">Activation through a belief or an emotion of an intermediate goal which is instrumental to the user's goal. It considers two possible applications: Rational Activation strategy or Emotional one.</td></tr>
<tr><td colspan="6">Induction of beliefs</td></tr>
<tr><td colspan="6">Argumentation about means-end implication. It represents the action-goal relation.</td></tr>
<tr><td>Appeal to Expert Opinion</td><td>Appeal to Popular Opinion</td><td>Appeal to Position to Know</td><td>Appeal to Friendly Personal Experience</td><td>Appeal to Examples</td><td>Others</td></tr>
</table>

In the *reasoning* phase, PORTIA exploits the information about the user in order to compute the degree of importance of the various -rational and emotional- goals, and infers the goals on which the persuasion strategy will focus. Using a "what-if" reasoning form it evaluates the persuasive power of different combination of strategies, and selects the most promising one, with respect to the goal of inducing in the user the intention to do a certain action.

In the *argumentation* phase, PORTIA constructs the arguments to express the strategy selected in the previous step by translating the output of the reasoning phase into a coherent discourse plan. The discourse plan is then translated by the dialogue manager into natural language messages used by the ECA as attempt to persuade the user.

8 Conclusions

The main goal of Intelligent Information Access is to provide a personalized and context-adapted access to information. In particular, systems implementing a conversational access to the information are enriched by exploiting human-computer interaction techniques (Mahmood et al. 2009). In this perspective, we present the architecture of an ECA which is able to exploit the knowledge conveyed by the user move in order to recognize her social attitude and goals and to behave 'believably' in its turn, by showing some forms of emotional intelligence. This research builds on prior work on affect modeling and dialog simulation. In particular, in this paper we combine social attitude and emotion modeling methods to build a scalable and open architecture for an emotionally and socially intelligent

ECA. In fact, each user move is rich of information (such as linguistic cues of social attitude) which goes beyond the pure content and meaning of sentences ('what user says'). These extra-rational information about the user state of mind can be exploited to enrich the user model and can be used by a socially and emotionally intelligent ECA, in order to tailor the dialogue strategy accordingly.

The two approaches to emotion and social attitude modeling have been validated in our previous research (Berry et al. 2005, de Rosis et al. 2007), with satisfying results.

We are aware of the limitations of our approach. In particular, translation of the user move meaning into symbolic form is currently implemented using a keyword-spotting based approach. In our future work, we plan to refine such analysis including contextual and acoustic information (Stolcke et al. 2000).

The main strength of the proposed ECA architecture is its openness and flexibility. In particular, we are able to simulate interactions in different conditions, by simply changing a few parameters describing the agent's personality. In this paper we show an example of adaptation by simulating the behavior of an empathic agent which reciprocates the social attitude of the user. In our future research we plan to perform evaluation studies in order to test which combination of personality traits of the agent best increases the user satisfaction. Moreover, we plan to investigate on the role that the interpersonal stances play in the display of emotions (De Carolis et al. 2001).

In our opinion, providing an intelligent and conversational access to information also requires the use of the most appropriate dialog strategy in order to increase the user intention to accept the system's suggestion. In this perspective, we describe an extension of our Agent's architecture with PORTIA, a module capable of simulating the persuasion process used by humans to convince someone to perform a given action.

Moreover, thanks to the independence of our architecture from the interaction mode, we plan to perform further investigation about spoken interaction. In particular, we will enrich the model for the analysis and interpretation of the user move using prosodic and acoustic parameters for improving the recognition of both (i) the actual communicative intention attached to the user move (De Carolis and Cozzolongo 2009) and (ii) the recognition of the user level of social attitude (de Rosis et al. 2007).

Acknowledgments. This research would not be without the encouragement, the suggestions and the teaching of Fiorella de Rosis. Her disappearance has been for us a sad loss, but her teaching incites us to continue the work.

References

Andersen, P.A., Guerrero, L.K.: Handbook of Communication and Emotions. In: Research, theory, applications and contexts. Academic Press, London (1998)

Berkovsky, S., Tsvi Kuflik, T., Ricci, F.: Cross-representation mediation of user models. User Modeling and User-Adapted Interaction 19(1-2), 35–63 (2009)

Berry, D.C., Butler, L.T., de Rosis, F.: Evaluating a realistic agent in an advice-giving task. International Journal of Man-Machine Studies 63(3), 304–327 (2005)

Bevacqua, E., Prepin, K., de Sevin, E., Niewiadomski, R., Pelachaud, C.: Reactive behaviors in SAIBA architecture. In: Workshop Towards a Standard Markup Language for Embodied Dialogue Acts, held in conjunction with AAMAS 2009 (2009)

Carbonell, J.C.: Towards a process model of human personality traits. Artificial Intelligence 15(1-2), 49–74 (1980)

Carofiglio, V., de Rosis, F., Novielli, N.: Dynamic User Modeling in Health Promotion Dialogs. In: Tao, J., Tan, T., Picard, R.W. (eds.) ACII 2005. LNCS, vol. 3784, pp. 723–730. Springer, Heidelberg (2005)

Carofiglio, V., de Rosis, F., Novielli, N.: Cognitive Emotion Modeling in Natural Language Communication. In: Tao, J., Tan, T. (eds.) Affective Information Processing, pp. 23–44. Springer, London (2008)

Carofiglio, V., de Rosis, F., Grassano, R.: Dynamic models of mixed emotion activation. In: Canamero, L., Aylett, R. (eds.) Animating expressive characters for social interactions, pp. 123–141. John Benjamins Publ. Co., Amsterdam (2008)

Cassell, J., Bickmore, T.: Negotiated collusion: modelling social language and its relationship effects in intelligent agents. User Modelling and User-Adapted Interaction 13(1-2), 89–132 (2003)

Clarizio, G., Mazzotta, I., Novielli, N., de Rosis, F.: Social Attitude Towards a Conversational Character. In: Proceedings of the 15th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2006, pp. 2–7 (2006)

Cooper, G.F., Herskovitz, E.: A Bayesian method for the induction of probabilistic networks from data. Machine Learning 9(4), 309–347 (1992)

De Carolis, B., Cozzolongo, G.: Interpretation of User’s Feedback in Human-Robot Interaction. Journal of Physical Agents 3(2), 47–58 (2009)

De Carolis, B.: MyTutor: A Personal Tutoring Agent. In: Panayiotopoulos, T., et al. (eds.) IVA 2005. LNCS (LNAI), vol. 3661, pp. 478–488. Springer, Heidelberg (2005)

De Carolis, B., Pelachaud, C., Poggi, I., de Rosis, F.: Behavior Planning for a Reflexive Agent. In: Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, IJCAI 2001, pp. 1059–1066 (2001)

De Carolis, B., Pelachaud, C., Poggi, I., Steedman, M.: APML, a Mark-up Language for Believable Behavior Generation. In: Prendinger, H., Ishizuka, M. (eds.) Life-like Characters, Tools, Affective Functions and Applications, pp. 65–85. Springer, Heidelberg (2004)

De Carolis, B.: Generating Mixed-Initiative Hypertexts: A Reactive Approach. In: IUI, pp. 71–78. ACM, New York (1999)

de Rosis, F., De Carolis, B., Carofiglio, V., Pizzutilo, S.: Shallow and inner forms of emotional intelligence in advisory dialog simulation. In: Prendinger, H., Ishizuka, M. (eds.) Life-like Characters, Tools, Affective Functions and Applications, pp. 271–294 (2003)

de Rosis, F., Batliner, A., Novielli, N., Steidl, S.: ‘You are Sooo Cool, Valentina!’ Recognizing Social Attitude in Speech-Based Dialogues with an ECA. In: Paiva, A.C.R., Prada, R., Picard, R.W. (eds.) ACII 2007. LNCS, vol. 4738, pp. 179–190. Springer, Heidelberg (2007)

de Rosis, F., Castelfranchi, C., Goldie, P., Carofiglio, V.: Cognitive Evaluations And Intuitive Appraisals: Can Emotion Models Handle Them Both? In: Draft Chapter of the HUMAINE Handbook. Springer, Heidelberg (in press) de Rosis, F., Novielli, N., Carofiglio, V., Cavalluzzi, A., De Carolis, B.: User Modeling And Adaptation In Health Promotion Dialogs With An Animated Character. International Journal of Biomedical Informatics 39(5), 514–531 (2006)

de Rosis, F., Pizzutilo, S., Russo, A., Berry, D.C., Molina, F.J.N.: Modeling the User Knowledge by Belief Networks. User Model and User-Adapted Interaction 2(4), 367–388 (1992)

Elliott, C., Siegle, G.: Variables influencing the intensity of simulated affective states. In: Proceedings of the AAAI Spring Symposium on Mental States 1993, pp. 58–67 (1993)

Hoorn, J.F., Konijn, E.A.: Perceiving and Experiencing Fictitional Characters: An integrative account. Japanese Psychological Research 45(4), 250–268 (2003)

Huang, H., Nishida, T., Cerekovic, A., Pandzic, I.S., Nakano, Y.: The design of a generic framework for integrating ECA components. In: Proceedings of the 7th international Joint Conference on Autonomous Agents and Multiagent Systems, International Conference on Autonomous Agents. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, Estoril, Portugal, May 12 - 16, vol. 1, pp. 128–135 (2008)

Jensen, F.V.: Bayesian Networks and Decision Graphs. Springer, Heidelberg (2001)

Kobsa, A.: User Modeling: Recent Work, Prospects and Hazards. In: Schneider-Hufschmidt, M., Kühme, T., Malinowski, U. (eds.) Adaptive User Interfaces: Principles and Practise. North Holland Elsevier, Amsterdam (1993)

Larsson, S., Traum, D.R.: Information state and dialogue management in the TRINDI dialogue move engine toolkit. Natural Language Engineering 6(3-4), 323–340 (2000)

Lee, C.M., Narayanan, S., Pieraccini, R.: Combining Acoustic and Language Information for Emotion Recognition. In: Proceedings of the seventh International Conference on Spoken Language Processing, ICSLP 2002 (2002)

Lee, E.J., Nass, C.: Effects of the form of representation and number of computer agents on conformity. In: Proceedings of Computer Human Interaction (CHI 1999), pp. 238–239 (1999)

Mahmood, T., Ricci, F., Venturini, A.: Learning Adaptive Reccomendation Strategies for Ondine Travel Planning. In: Information and Communication Technologies in Tourism 2009, pp. 149–160. Springer, Wien (2009)

Mazzotta, I., de Rosis, F.: Artifices for persuading to improve eating habits. In: AAAI Spring Symposium on 'Argumentation for consumers of health care', pp. 76–85 (2006); Technical Report SS-06-01

Mazzotta, I., de Rosis, F., Carofiglio, V.: PORTIA: a user-adapted persuasion system in the healthy eating domain. IEEE Intelligent Systems 22(6), 42–51 (2007)

Mazzotta, I., Silvestri, V., de Rosis, F.: Emotional And Non Emotional Persuasion Strength. In: Proceedings of AISB 2008, Symposium on 'Persuasive Technology', pp. 14–21 (2008)

Miceli, M., de Rosis, F., Poggi, I.: Emotional and non-emotional persuasion. Applied Artificial Intelligence: an International Journal 20(10), 849–880 (2006)

Neumann, B.C.: Scale in Distributed Systems. In: Casavant, T., Singhal, M. (eds.) Readings in Distributed Computing Systems, pp. 463–489. IEEE Computer Society Press, Los Alamitos (1994)

Novielli, N., de Rosis, F., Mazzotta, I.: User attitude towards an embodied conversational agent: Effects of the interaction mode. Journal of Pragmatics, doi:10.1016/j.pragma.2009.12.016 (2010)

Oatley, K., Johnson-Laird, P.N.: Towards a Cognitive Theory of Emotions. Cognition and Emotion 1, 29–50 (1987)

Ortony, A., Clore, G.L., Collins, A.: The cognitive structure of emotions. Cambridge University Press, Cambridge (1988)
Paiva, A. (ed.): Empathic Agents. Workshop in conjunction with AAMAS 2004 (2004)
Pearl, J.: Probabilistic Reasoning in Expert Systems: Networks of Plausible Reasoning. Morgan Kaufmann, San Mateo (1988)
Petty, R.E., Cacioppo, J.T.: The elaboration likelihood model of persuasion. In: Berkowitz, L. (ed.) Advances in experimental social psychology, vol. 19, pp. 123–205. Academic Press, New York (1986)
Picard, R.W.: Affective Computing. MIT Press, Cambridge (1997)
Poggi, I., Pelachaud, C., de Rosis, F.: Eye Communication in a Conversational 3D Synthetic Agent. AI Commun. 13(3), 169–182 (2000)
Polhemus, L., Shih, L.F., Swan, K.: Virtual interactivity: the representation of social presence in an on line discussion. In: Annual Meeting of the American Educational Research Association (2001)
Reeves, B., Nass, C.: The Media Equation: How people treat computers, television and new media like real people and places. Cambridge University Press, New York (1996)
Schröder, M. (2010) The SEMAINE API: Towards a Standards-Based Framework for Building Emotion-Oriented Systems. Advances in Human-Computer Interaction. Article ID 319406, 21 pages (2010) doi:10.1155/2010/319406
Stolcke, A., Coccaro, N., Bates, R., Taylor, P., Van Ess-Dykema, C., Ries, K., Shriberg, E., Jurafsky, D., Martin, R., Meteer, M.: Dialogue act modeling for automatic tagging and recognition of conversational speech. Computational Linguistics 26(3), 339–373 (2000)
Swan, K.: Immediacy, social presence and asynchronous discussion. In: Bourne, J., Moore, J.C. (eds.) Elements of quality online education, Nedham, MA. Sloan Center For Online Education, vol. 3, pp. 157–172 (2002)
Vilhjálmsson, H., Cantelmo, N., Cassell, J., Chafai, N.E., Kipp, M., Kopp, S., Mancini, M., Marsella, S., Marshall, A.N., Pelachaud, C., Ruttkay, Z., Thórisson, K.R., Welbergen, H., Werf, R.J.: The Behavior Markup Language: Recent Developments and Challenges. In: Pelachaud, C., Martin, J.-C., André, E., Chollet, G., Karpouzis, K., Pelé, D. (eds.) IVA 2007. LNCS (LNAI), vol. 4722, pp. 99–111. Springer, Heidelberg (2007)
Walton, D.N.: Argumentation Schemes for Presumptive Reasoning. N.J. Erlbaum, Mahwah (1996)
Walton, D.N.: The place of emotion in argument. The Pennsylvania State University Press (1992)
Walton, D.N.: What is reasoning? What is an argument? Journal of Philosophy 87, 399–419 (1990)