Cooperating with Personality-Rich Interface Agents

Sebastiano Pizzutilo, Berardina De Carolis and Fiorella de Rosis

Department of Informatics, University of Bari

{ pizzutilo, decarolis, derosis}@di.uniba.it

Abstract

Animated Agents are endowed with personality (and emotions), with the aim of increasing their believability and of establishing an empathic relationship with the User: the Five Factor Model is the reference schema most frequently employed to this aim. In this paper, we claim that, to endow Agents with 'social intelligence', these 'communication' traits should be integrated with 'cooperation' attitudes; we describe our experience in building an Agent that combines the two personality aspects and discuss the problems still open.

Why Personality-Rich Interface Agents

In the near future, computers will either 'disappear', to ubiquitously pervade life environment in a not immediately perceivable way, or will tend to take the external appearance and the internal behaviour of a human being, to undertake a friendship relation with the User. The two scenarios will probably coexist, to apply in different contexts. In both cases, endowing agents with some form of social intelligence appears to be a crucial need: if reasoning, help and control abilities are distributed among specialised agents integrated with objects of daily life, some form of communication and cooperation among these agents is needed, to avoid conflicting goals and behaviours. On the other side, an embodied agent that tries to give the Users the illusion of 'cooperating with a partner or a friend' should be able to understand them, to help them in solving problems, to find ways of coming to a mediated solution in case of conflicts, and so on. To make easier, to the User, foreseeing how the Agents will behave, they should harmonise their external appearance with their internal behaviour; they should understand how to adapt to the User needs and moods and should, finally, enable the Users to 'select a different partner', when the Agent with which they interact does not fit with their preferences.

If short-term variations in behaviour and appearance of Embodied Agents have been metaphorically represented in terms of '*emotional states*', more stable differences have been represented in terms of '*personality traits*'. Endowing Socially Intelligent Agents with a personality involves solving the following problem issues*:*

1. *which forms of social intelligence* these Agents should have and how they may be translated in terms of personality traits;
2. *how a trait may be represented* in the Agent's 'mind', that is in its mental state and reasoning style;
3. *how various traits may be combined* in the same individual and, finally,
4. *how one or more traits may manifest* in the external behaviour of the Agent.

In this paper, we discuss our experience in building an Interface Agent that cooperates with the User in performing the tasks included in a software application; the Agent is designed and built with a BDI architecture: we will focus our description on the way that we formalised its cooperation attitude.

Which Dimensions of Personality

Research on personality-based HCI has be driven by results of studies about human intelligence: in particular, the 'Five Factor Model' (FFM) and the 'Interpersonal Circumplex Model' (IC).

The FFM is a response that psychologists found about 20 years ago to the need of defining '*the most important ways in which individuals differ in their enduring emotional, interpersonal, experiential, attitudinal and motivational styles*' (McCrae and John, 1992). The five dimensions[1] are an interpretation of results of applying factor analysis to questionnaires submitted to various groups of subjects, by researchers of different groups; their meaning is a subjective interpretation of the set of variables they 'explain', and is described with natural language terms. 'Sociability' or 'Social closeness' is associated, in particular, with Extraversion. Other authors suggest that 'socialisation might be a higher order factor combining high Agreableness and Conscientiousness with low Neuroticism' (McCrae and John,1992).

The second method employed to categorise human personalities is Wiggin's measure of IC, whose axes are 'Dominance' and 'Affiliation'. Whether the two factorisation criteria are related is not fully clear: some authors identify Extraversion with Dominance, while others argue that Extraversion is best seen as located midway between Dominance and Warmth (McCrae and John, 1992).

Researchers in HCI employed the two mentioned factorisation criteria to enrich interfaces with a personality. Some notable examples: Nass and colleagues studied graphical interfaces in terms of Dominance (Nass et al, 1995) and Agent-based interfaces in terms of Extraversion

[1] Extraversion, Agreableness, Conscientiousness, Neuroticism and Openness.

(Nass et al, 2000); Dryer (1998) adopted the IC model; André et al (2000) attach Extraversion and Agreableness to the members of their 'Presentation Teams'; Ball and Breese (2000) include Dominance and Friendliness in their modeling of personality-related observable behavior. Taylor and colleagues (1998) studied how the FFM traits can be portrayed using voice and gestures.

The advantage of referring to the two mentioned models, for computer scientists, is to have a largely accepted frame of reference, with various useful fall-outs, such as a definition of the way that every personality factor manifests itself in the external behavior. The main disadvantage, in our view, is that these personality traits refer to a characterisation of 'affective expression' and 'communication style' rather than to mental social attitudes. They are therefore very useful for endowing Agents with a 'pleasant' and 'believable' appearance, but not to express diversification in social relationships.

Another difficulty in employing the cited models is that traits are defined through natural language descriptions: as a consequence, they are not particularly suited to be formalised into the 'mental state' of a BDI agent[2]. Building 'socially intelligent agents' requires a theory that enables combining personality-driven reasoning attitudes with a characterisation of the communication style: social psychologists should, if possible, collaborate with computer scientists in developing these theories and formalisms. The first and most relevant contribution in this direction was due to Carbonell (1980), who formalised personalities in terms of combinations of degrees of importance assigned to goals, and to Cohen and Leveque (1990), who saw them as dichotomic attributes that trigger reasoning rules (for instance, 'sincere' or 'helpful'). A third example, to which we will refer in particular in this paper, is Castelfranchi and Falcone's (1998) theory of cooperation in multi-agent systems.

Cooperation Personalities

Although emotional and affective expressions may contribute to increase the Agents' friendliness, it is the level of help they provide to the User, their 'cooperation attitude', that really drives their acceptability. The *level of help* an Agent provides should not be equal to all users, but should be tailored to their needs and attitudes towards computers in general and towards the specific software to which they are applied in particular; these needs and attitudes may be synthetized in *a level of delegation of tasks* that the User adopts towards the Agent. To select the helping attitude that best suits to the User needs, the Agent has to be endowed with a 'mental state' and with a reasoning ability, that enable it to observe the User, to represent her expected abilities and needs in a User Model and to plan the 'best' response in every context, according to the content of this Model.

To be represented into the Agent's (or the User's) 'mental state', the definition of 'cooperation attitude' should be formalised with some logical language. As we said, we applied, in our research on Animated Agents, the theory of cooperation among agents that was defined by Castelfranchi and Falcone (1998). We had already applied this theory to formalise the mental state of agents and their reasoning abilities in a previous Project (GOLEM), where two agents with different cooperation attitudes interacted to solve a problem in a toy domain, in situations of possible conflict (Castelfranchi et al, 1998). This enabled us to investigate the effect of combining different delegation and help personality traits into the same Agent and of enabling Agents with different personalities to encounter, by observing the effects of these combinations. Our Agents were not, in that case, embodied with a life-like appearance: with XDM-Agent, we extended our research on cooperation attitudes in the direction of Embodied Animated Agents.

XDM-Agent

XDM-Agent[3] is an embodied animated character that cooperates with the User in performing the tasks included in a given application: its cooperation attitude changes according to the User and the context and its external appearance is programmed to agree with its communication trait. Although XDM-Agent is domain-independent, in this paper we take electronic mail as a case study, to show some examples of how it behaves in guiding the use of Eudora.

The present version of Eudora is not very generous in providing help to the User. In a software of large use like this, writing and sending a letter, checking mail and arranging email in files should be very natural: the system should therefore cooperate with Users in trying to make these (and other) tasks as easy to perform as possible. The first goal of XDM-Agent is then "*to make sure that the User performs the main tasks included in Eudora without too much effort*". At the same time, the Agent should avoid providing too much help when this is not needed or when the User prefers to be left alone in making her experience. A second goal is therefore "*to make sure that the User does not see the Agent as too much intrusive, or ennoying*".

These general goals may specialise into more specific ones, according to the 'cooperation attitude' of the Agent. In deciding the level and the type of help to provide, XDM-Agent should consider, at the same time, the User

[2] that is, of an Agent whose reasoning ability is implemented through a processing of 'Belief, Desire and Intentions' (Rao and Georgeff, 1991)

[3] XDM-Agent was originally designed to help the User in understanding how a given software application may be used: it is therefore an Animated User Manual that employes, in generating explanations, a model of the application represented with a formalism that we called XDM (Context-Sensitive Dialogue Modeling: see de Rosis et al, 1998 and de Rosis et al, 2000).

experience and her 'delegation attitude': we remember, in **Table 1,** how delegation and help attitudes have been defined by Castelfranchi and Falcone.

Delegation attitudes
a ***lazy*** agent always delegates tasks if thereis another agent that is able to take care of them; it acts by itself only when there is no alternative;
a ***hanger-on*** tends to never act by itself;
a ***delegating-if-needed*** asks for help only if it is not able to do the task by itself;
a ***never-delegating*** considers that tasks should only be achieved if it canperform them.
Helping attitudes
a ***hyper-cooperative*** always helps if he can
a ***benevolent*** first checks that the other agent could not do the action by itself
a ***supplier*** first checks that the request does not conflict with its own goals
a ***selfish*** helps only when the requested action fits with its own goals
a ***non-helper*** never helps, on principle
Helping levels
a ***literal helper*** restricts itself to considering the requested action
a ***overhelper*** goes beyond this request, to hypothesize a delegating agent's higher order goal, and helps accordingly
a ***subhelper*** performs only a subset of the requested plan
a ***critical helper*** modifies the delegated plan by, at the same time, considering literally the request or going behind it or responding only partially to it

Table 1: some delegation and help traits and levels (from Castelfranchi et al, 1998).

If formalised into personality-related decision rules, the definitions in this Table show that the Agent's decision of whether and how to help the User relies on the following knowledge sources:

Own Mental State:

the Agent's goals (Goal XDM (T g)) and abilities (Bel XDM (CanDo XDM a)), and the actions it intends to perform (Bel XDM (IntToDo XDM a)) should be represented explicitly in this KB;

Domain Knowledge:

given a domain goal, XDM should know all the plans that enable achieving this goal, in the application:

∀g ∀p (Domain-Goal g) ∧ (Domain-Plan p) ∧ (Achieves p g) ⇒ (KnowAbout XDM g) ∧ (KnowAbout XDM p) ∧ (Know XDM (Achieves p g)).

It should know, as well, the individual steps of every domain-plan:

∀p ∀a (Domain-Plan p) ∧ (Domain-action a) ∧ (Step a p)⇒ (KnowAbout XDM p) ∧ (KnowAbout XDM a) ∧ (Know XDM (Step a p)).

These rules establish that XDM has a 'complete' knowledge of the tasks the application enables performing and how each of them may be performed. This knowledge has to be represented with some formalism: we employ Coloured Petri Nets to represent the interface of user-adapted applications, with the tasks the application enables performing, the order in which these tasks have to be performed and their relation with interface objects (de Rosis et al, 1998).

User Model:

the agent should have some hypothesis about:

- the User goals, both in general and in specific phases of interaction:
  ∀ g (Goal U (T g)) ⇒ (Bel XDM (Goal U (T g))),
- her abilities:
  ∀ a (CanDo U a) ⇒ (Bel XDM (CanDo U a)) and
- what the User expects the Agent to do, in every phase of interaction:
  ∀ a (Goal U (IntToDo XDM a)) ⇒
  (Bel XDM Goal U (IntToDo XDM a)).

This may be a default, stereotypical knowledge about the User that is settled at the beginning of interaction: the stereotype is related to the delegation attitude of the User and to his/her level of experience in the domain to which the application refers. Ideally, the model should be updated dynamically, according to a plan-recognition activity performed by the Agent.

Reasoning Rules:

the Agent employs this knowledge to take decisions about the level of help to provide in any phase of interaction, according to its 'helping attitude', which is represented as a set of 'reasoning rules'. As we show in Table 1, if, for instance, XDM-Agent is a *benevolent*, it will respond to the User's (implicit or explicit) requests of performing actions that it presumes the User is not able to do:

Rule R1

∀a [(Bel XDM (Goal U (IntToDo XDM a))) ∧ (Bel XDM ¬ (CanDo U a))))) ∧ (Bel XDM (CanDo XDM a))] ⇒ (Bel XDM (IntToDo XDM a))

If, on the contrary, the Agent is a *supplier*, it will do the requested action only if this does not conflict with its own goals:

Rule R2

∀a [(Bel XDM (Goal U (IntToDo XDM a))) ∧ (Bel XDM (CanDo XDM a)) ∧ (¬∃ g (Goal XDM (T g)) ∧ (Bel XDM (Conflicts a g))] ⇒ (Bel XDM (IntToDo XDM a))

... and so on for the other personality traits.

An Example

Let us consider the domain-goal g = "to specify the Recipient's address in a e-mail', and let us assume that this goal may be achieved through several plans, involving the following elementary steps (actions):

a1: to parse the address string
a2: to check whether the string is correct
a3 to display a generic error message
a4: to specify which error the string includes
a5: to ask the User to correct the string
a6: to propose alternatives on how to correct the string
a7: to look for a similar string in the Address Book
a8: to check whether the similar address found in the Address Book corresponds to the intended Recipient's address

a9: to substitute the similar address found in the Address Book in the place of the erroneous one,
a10: to validate whether the User agrees with the substitution.

Alternative correct plans for achieving g when the address string is not correct are the following:

p1 = (a1, a2, a3), (generic error message)
p2 = (a1, a2, a4, a5), (detailed error message)
p3 = (a1, a2, a4, a6, a5), (suggestion of correction)
p4 = (a1, a2, a7, a8, a9, a10) (automatic correction)

,... (and maybe still others). These plans correspond to increasing levels of cooperation that XDM might adopt in helping the User to perform this task.

Let us assume that our Agent is a *benevolent*: in deciding whether to help the User in reaching a domain-goal g it will have to check, first of all, whether (Bel XDM (Goal U (T g)). We said that the Agent knows about several plans bringing to g and, for each of them, knows how to perform all the steps included in the plan. Let us assume that no conflict exists between g and the Agent's goals. By applying rule R1, XDM will come to the decision to do its best to help the user in reaching the goal g, by directly performing all the steps of the plan: for instance, if XDM selects the plan p2, it will perform, in sequence, actions a1, a2, a4, a5. As we said, however, this plan is not unique: the Agent might select other alternatives, for instance p1, p3, p4: each of these alternatives corresponds to a different *level of help* that the Agent will provide to the User.

The level of help the Agent is willing to provide may be seen, as well, as a personality trait (see, again, Table 1). If, for instance, XDM-Agent is a *literal helper*, which interprets literally the User need to avoid errors in the Recipient's address, it will apply the plan p2. If, on the contrary, the Agent is a *overhelper*, that goes beyond the User request of help to hypothesize her higher order goal (for instance, to be helped in correcting the address, if possible), it will apply plan p3 or even p4. A *subhelper* will apply plan p1, by only sending a generic error message: this is what Eudora does at present if the User tries to send a message without specifying any address. If, finally, the User asks XDM-Agent to suggest how to correct the string (plan p3) and the Agent is not able to perform action a6 (while it knows how to perform actions a7, a8) and it is a *critical helper*, it will select and apply, instead, plan p4.

How to Combine Personality Traits

In multiagent cooperation, an Agent may find itself in the position of delegating some task or helping other Agents, in different phases of a complex activity. A theory is therefore needed to establish how delegation and helping attitudes may combine in the same Agent. In XDM-Agent, cooperation with the User is very 'altruistic', as the User is not expected to help the Agent in performing a task: we had not to define, consequently, how the two personality traits are combined. Some general thoughts about this topic may be found in (Castelfranchi et al, 1998).

The Agent's reasoning on whether to help the User in performing a task ends up with an intentional state (to perform an individual action, an entire plan, part of a plan etc). This intentional state is transformed into action, that may include communication with the User: for instance, an *overhelper* Agent which decides to perform plan p3 will interact with the User to specify the error included in the string, to propose alternatives on how the string might be corrected and to ask her to correct it. In this phase, the Agent will adopt a communication personality trait: for instance, it might do it in an 'extroverted' or an 'introverted' way. The question then is: *how should cooperation and communication personalities be combined?* Is it more reasonable to assume that a hypercooperative agent is an extroverted or an introverted? We do not have, at present, an answer to this question. In the present prototype, we implemented only two personalities (a *benevolent* and a *supplier*) and we associated the 'benevolent' trait with the 'extroverted' one, the 'supplier' with the 'introverted'.

How to Match Agent's and User's Personalities

The User desire to receive help may be formalised, as well, in personality terms: if the User is a '*lazy*', she expects to receive, from XDM, some cooperation in completing a task, even if she would be able to do it by herself (and therefore, irrespectively of her 'level of experience').

Rule R3

$\forall a\, \forall g\, [((Goal\ U\ (T\ g)) \wedge (Bel\ U\ (Achieves\ a\ g)) \wedge (Bel\ U\ (CanDo\ XDM\ a))) \Rightarrow (Goal\ U\ (IntToDo\ XDM\ a))\ ]$

If, on the contrary, the User is a '*delegating-if-needed*', she will expect to receive some help only if she is not able to do the job by herself (for instance, if she is a 'novice' in sending emails):

Rule R4

$\forall a\, \forall g\, [((Goal\ U\ (T\ g)) \wedge (Bel\ U\ (Achieves\ a\ g)) \wedge (Bel\ U\ \neg(CanDo\ U\ a)) \wedge (Bel\ U\ (CanDo\ XDM\ a))) \Rightarrow (Goal\ U\ (IntToDo\ XDM\ a))\ ].$

Providing help in this type of job to an 'expert' and 'delegating-if-needed' User will be seen as a kind of intrusiveness, that will violate the Agent's goal to 'avoid annoying the User'.

In our first prototype of XDM-Agent, the Agent's cooperation personality (and therefore its helping behaviour) may be settled by the User at the beginning of interaction or may be selected according to some hypothesis about the User. As we said before, ideally the Agent should be endowed with a plan recognition ability, that enables it to update dynamically its image of the User personality and abilities: notice that, while recognising communication traits requires observing the 'external' (verbal and nonverbal) behavior of the User, inferring her cooperation attitude requires reasoning on the history of

interaction (a 'cognitive diagnosis' task that we studied, in probabilistic terms, in (de Rosis et al, in press)).

Once some hypothesis about the User's delegation personality exists, how should the Agent's helping personality be settled?. One of the controversial results of research about communication personalities in HCI is whether the similarity or the complementarity principles hold: that is, whether an 'extroverted' interface should be proposed to an 'extroverted' User, or the contrary, and so on. When cooperation personalities are considered, the question becomes the following: "How much should an Interface Agent help a User? How much importance should be given to the User experience (and therefore her abilities in performing a given task), how much to her propensity to delegate that task?". In our opinion, the answer to this question is not unique. If XDM-Agent's goals are those mentioned before, that is *"to make sure that the User performs the main tasks included in Eudora without too much effort*" and "*to make sure that the User does not see the Agent as too much intrusive, or ennoying*", then the following combination rules may be adopted:

- (DelegatingIfNeeded U) ⇒ (Benevolent XDM):
  XDM-Agent helps delegating-if-needed Users only if it presumes that they cannot do the action by themselves;
- (Lazy U) ⇒ (Supplier XDM):
  XDM-Agent does its best to help lazy users, unless this conflicts with its own goals;

.... and so on.

However, if the Agent has also the goal *"to make sure that Users exercise their abilities*" (such as in Tutoring Systems), then the matching criteria will be different; for instance:

- (Lazy U) ⇒ (Benevolent XDM):
  XDM-Agent helps a lazy User only after checking that she is not able to do the job by herself. In this case, the Agent's cooperation behavior will be combined with a communication behavior (for instance, Agreableness) that warmly encourages the User in trying to solve the problem by herself.

How to Manifest the Agent's Personality

XDM-Agent was built with the MS-Agent development software, as a character which can take several 'bodies', can move on the display to indicate objects, can make several arm and body gestures, can speak and write a short text on a balloon. To enable the User to foresee how the Agent will behave in a given circumstance and to insure that its external appearance (its 'Body') is, in a way, consistent with its internal behaviour (its 'Mind'), the ideal would be to match the Agent's appearance with its helping personality: however, as we said, no data are available on how cooperation traits manifest themselves, while literature is rich on how communication traits are externalised into verbal and nonverbal behavior. At present, therefore, XDM-Agent's body only depends on its communication personality (that is, on whether it is an extroverted or an introverted). To enhance this difference, we associate a different character with each of them (*Genie* with the benevolent&extroverted and *Robby* with the supplier&introverted). **Table 2** summarises the main features of the two characters.

Robby	**Genie**
supplier	benevolent
introverted	extroverted
is rather 'passive': says the minimum and waits for the user's request of explanation	is very 'active': takes the initiative and provides detailed explanations
employs 'light' linguistic expressions, with indirect and uncertain phrasing (suggestions)	employs 'strong' linguistic expressions, with direct and confident phrasing (commands)
gestures the minimum: minumum locomotion, limited movements of arms and body avoids getting close to the user	gestures are more 'expansive': more locomotion, wider movements of arms and body gets close to the user when needed
speaks slow	speaks high

Table 2: Relationship between XDM-Agent's personality traits and their external behavior.

The MS-Agent technology we employed enabled us to program the Agent to perform a minimal part of the gestures we would need: we are therefore working, at the same time, at a more refined Animated Agent that can adapt its face, mouth and gaze to its high-level goals, beliefs and emotional states: this enables us to directly link individual components of the Agent's mental state to its verbal and non-verbal behaviour, through a set of personality-related activation rules. In this way, we plan to obtain a much more refined matching between the Agent's Mind and its Body (Poggi et al, in press).

Conclusion

Animated Agents tend to be endowed with a personality and with the possibility to feel and to display emotions, for several reasons. Some examples: in Tutoring Systems, display of emotions enables the Agent to show to the students that it cares about them and that it is sensitive to their emotions; it favours convey of enthusiasm and contributes to insure that the student enjoys during learning (Elliott et al, 1997). In Information- Providing Systems, personality traits contribute to specify a motivational profile of the Agent and to orient the dialog accordingly (André et al, 2000). Personality and emotions are attached to Personal Service Assistants to better 'anthropomorphize' them (Arafa et al, 1998). As we said at the beginning of this paper, personality traits that are attached to the Agents, in these well-known Projects,

reproduce the Big-Five Factors that seem to characterise social relations in humans: although a 'decisional level' (for instance, high-level planning, in André et al) is usually separated from a 'communication level' (realisation of verbal and nonverbal behaviours), both levels are driven by the same personality traits.

Among the traits that have been considered so far, 'Dominance/Submissiveness' is the only one that in a way relates to the definition of cooperation attitudes: Nass and colleagues' definition for the two extremes of this trait includes a mixture of help and delegation attitudes:

"Dominant individuals tend to try to exercise power over the minds or behaviour of others,...to make decisions for others and direct others to take certain actions.... They are marked by the following verb-phrase descriptions:

(1) able to give orders,
(2) talks others into doing what he/she wants;
(3) often assumes responsibility.

Conversely, submissiveness is behaviourally marked by the following verb-phrase descriptions:

(1) easy led
(2) lets others make decisions
(3) avoids responsibility" (Nass et al, 1995)

'Dominants' are therefore those who pretend that others help them when they need it; at the same time, they tend to help others by assuming responsibilities on themselves. 'Submissive', on the contrary, tend, on one side, to obey to orders and, on the other side, to delegate actions and responsibilities whenever possible.

This model seems, however, to consider only some combinations of cooperation and communication attitudes that need, on the contrary, to be studied and modeled separately and more indepth. We claim that Castelfranchi and Falcone's theory of cooperation might contribute to such a goal, and the first results obtained with the XDM-Agent prototype encourages us to go on in this direction. As we said, however, much work has still to be done to understand how psychologically plausible configurations of traits may be defined, how they evolve dynamically during interaction and how they are externalised.

References

André, E., Rist, T:, van Mulken, S., Klesen, M. and Baldes, S. 2000. The automated design of believable dialogues for animated presentation teams. In Cassel, J, Sullivan, J., Prevost, S. and Churchill, E. (Eds): *Embodied Conversational Agents*. The MIT Press.

Arafa, J., Charlton, P., Mamdani, A. and Fehin, P. 1998. Designing and building Personal Service Assistants with personality. In , Prevost, S. and Churchill, E. (Eds). *Proceedings of the Workshop on Embodied Conversational Characters*.

Ball, G. and Breese, J. 2000. Emotion and Personality in a Conversational Agent. In Cassel, J, Sullivan, J., Prevost, S. and Churchill, E. (Eds): *Embodied Conversational Agents*. The MIT Press.

Carbonell, J. 1980. Towards a process model of human personality traits. *Artificial Intelligence*, 15.

Castelfranchi, C. and Falcone, R. 1998. Towards a theory of delegation for agent-based systems. *Robotics and Autonomous Systems, Special Issue on Multi-Agent Rationality*.

Castelfranchi, C., de Rosis, F., Falcone, R. and Pizzutilo, S. 1998. Personality traits and social attitudes in multiagent cooperation. *Applied Artificial Intelligence,* 12: 7-8.

Cohen, P. R. and Levesque, H. 1990. Rational interaction as the basis for communication. In Cohen, P. R., Morgan, J. and Pollack, M. E. (Eds). *Intentions in Communication*. The MIT Press.

de Rosis, F., Pizzutilo, S. and De Carolis, B. 1998. Formal description and evaluation of user-adapted interfaces. *International Journal of Human-Computer Studies*.

de Rosis, F., De Carolis, B and Pizzutilo, S. 2000. Automated Generation of Agent Behaviour from Formal Models of Interaction. *Proceedings of the Conference 'Advanced Visual Interfaces', 2000*.

de Rosis, F., Covino, E., Falcone, R, and Castelfranchi, C. in press. Bayesian cognitive diagnosis in believable multiagent systems. In press. In: Williams, M. A. and Rott, H. (Eds), *Frontiers of Belief Revision*. 'Applied Logic' Series of Kluwer.

Dryer D.C. 1998. Dominance and valence: a two-factor model for emotion in HCI. *Proceedings of AAAI 1998*.

Elliott, C., Lester, J.C. and Rickel, J. 1997. Interpreting affective computing into animated tutoring agents. *Workshop on Intelligent Interface Agents,* IJCAI 97.

McCrae, R. R. and John, O.1992. An introduction to the five-factor model and its applications. *Journal of Personality,* 60.

Nass, C., Moon, Y., Fogg, B. J., Reeves, B. and Dryer, D.C. 1995. Can computer personalities be human personalities? *International Journal of Human-Computer Studies*, 43.

Nass, C., Isbister, K. and Lee, E-J. 2000. Truth is beauty: Researching Embodied Conversational Agents. In Cassel, J, Sullivan, J., Prevost, S. and Churchill, E. (Eds): *Embodied Conversational Agents*. The MIT Press.

Poggi, I., Pélachaud, C. and de Rosis, F. Eye communication in a conversational 3D synthetic agent. *AI Communications*, in press.

Rao, A. S: and Georgeff, M. P. 1991. Modeling rational agents within a BDI architecture. In *Principles of knowledge representation and reasoning*.

Taylor, I.C., McInness, F. R., Love, S., Foster, J.C. and Jack, M. A. 1998. Providing animated characters with designated personality profiles. In *Proceedings of the Workshop on Embodied Conversational Characters*, Prevost, S. and Churchill, E. (eds).