Agents That Talk Back (Sometimes): Filter Programs for Affective Communication

Helmut Prendinger and Mitsuru Ishizuka
Department of Information and Communication Engineering
School of Engineering, University of Tokyo
7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan
{helmut,ishizuka}@miv.t.u-tokyo.ac.jp

Abstract

This paper introduces a model of interaction between users and animated agents as well as inter-agent interaction that supports basic features of affective conversation. As essential requirements for animated agents' capability to engage in and exhibit affective communication we motivate reasoning about emotion and emotion expression, personality, and social role awareness. The main contribution of our paper is the discussion of so-called 'filter programs' that may qualify an agent's expression of its emotional state by its personality and the social context. All of the mental concepts that determine emotion expression, such as emotional state, personality, standards, and attitudes, have associated intensities for fine-tuning the agent's reactions in user-adapted environments.

Introduction

Recent years show a growing interest in animated agents as conversational partners for a variety of tasks that are typically performed by humans. Animated agents are used

- as virtual *tutors* in interactive learning environments (e.g., de Rosis *et al.* 1999, Johnson *et al.* 2000),
- as virtual *presenters* on the web (e.g., André *et al.* 2000, Ishizuka *et al.* 2000),
- and as virtual *actors* for entertainment (e.g., Rousseau and Hayes-Roth 1998, Paiva *et al.* 2001) and language conversation training (e.g., Prendinger and Ishizuka 2001c; 2001b).

There is general agreement that emotion and personality are key ingredients for pedagogically effective or dramatically interesting, in short, *believable* characters. However, besides the diversity of psychological models of emotion and personality, there is no consensus about which other mental concepts should be modelled to achieve believable behavior. Allen (1999) considers a broad range of 'higher-level' mental concepts—personality, attitudes, standards, moods, emotions, desires, intentions, and plans—which he calls (motivational) *control states*. A mental concept is considered as a control state if it might function as a predictor of behavior. If we can say "she does this because she is in a bad mood", without referring to other of the observed agent's mental concepts (e.g., personality), it is a good indicator that mood is a control state. The Émile system (Gratch 2000), for instance, successfully integrates reasoning about emotion and reasoning about plans. Our current model considers the following mental concepts (or control states): emotion, personality, mood, attitudes, standards, goals, and a control state that we call 'social role awareness'. The integration and interaction of those concepts allows for a broad variety of believable agent behaviors.

Our goal is to develop a general framework for affective communication that covers various forms of user-agent and inter-agent interactions. Specifically, we will propose components of agents' mental models that enable them to process emotions and show affective behavior. A salient feature of our model is the distinction between reasoning about emotion (affective reasoning) and emotion expression. While affective reasoning involves reasoning about events and an agent's goals, standards, and attitudes, so-called *filter programs* qualify the expression of the inferred emotional state by the agent's personality (or mood) and the social context. Important issues include the integration of the intensities of the various control states to the (overall) intensity of emotion expression, as well as the impact of the user's (or another agent's) communicative act on the response of the agent.

Our approach is used to improve English conversation skills of native speakers of Japanese, where interactions are set up as role-playing dramas and games. The programmable interface of the Microsoft Agent package (Microsoft 1998) is used to run our interactive role-playing scenarios. Although these off-the-shelf agent characters are quite restricted in the number of behaviors, the package comes ready with a speech recognizer and text-to-speech engine that allow client-side execution in a web browser.

The rest of the paper is organized as follows. In the next section, we sketch our communication model. After that, we describe an affective reasoner for reasoning about emotion, and report on an influential theory of emotion expression. The following section introduces so-called 'filter programs' that function as a 'filter' be-

tween the agent's affective state and emotion expression. First, we briefly describe a simple model of personality and social role awareness, and then give some examples of filter rules. We also summarize our findings from an empirical study on users' reactions to socially aware (and unaware) animated agents. In the final section, we discuss and conclude the paper.

Simulating Conversations

We employ Moulin and Rousseau's (2000) approach to model and simulate conversations, which provides a rich framework for many aspects of inter-agent communication. In a conversation, multiple (locutor-)agents communicate through various channels, such as verbal utterances, gestures, body movement, and facial display. Following (Moulin & Rousseau 2000), we distinguish three levels of communication. Communication maintenance and turn-taking is performed at the *communication* level. At the *conceptual* level, agents transfer concepts (mental states). Finally, at the *social* level, agents manage and respect social relationships that hold between agents. In the present paper, we will focus on the conceptual and the social level. The communication level is described, e.g., in (Cassell & Thórisson 1999).

At the conceptual level, information is exchanged between agents as simplified symbolic representations of the utterance. If an agent orders coffee, this is simply represented as *order_coffee*. The verbal act corresponding to "order coffee", typically accompanied by non-verbal acts, is determined by the social context of the conversation and the agent's characteristics (e.g., personality).

Consider an agent called 'Al' ordering coffee from an agent called 'James', by saying "I would like to order coffee". As a basic interaction protocol between agents, we propose communicative acts of the form

com_act(al,james,order_coffee,polite,happiness,s0)

where 'al' is the speaker, 'james' is the hearer, 'order_coffee' is the conveyed information, 'polite' is a qualitative evaluation of the linguistic style of the utterance, 'happiness' refers to Al's emotion expression (e.g., smiling), and *s0* denotes the situation in which the utterance takes place. As an emotion (expression) recognition module is not part of our system, 'neutral' is set as the default value for users' emotion expression.

Each agent involved in a conversation is assumed to have its own mental model. A mental model contains different kinds of entities (components), including world knowledge (beliefs), and representations of higher-order mental concepts (emotions, personality traits, standards, attitudes, goals). In the following sections, we will discuss control programs that process those components: an affective reasoning program for inferring an agent's emotional state, and filter programs for reasoning about personality and the social context. Furthermore, we will discuss the output—an agent's emotion expression—which is the expression of its affective state qualified by its personality and the social setting in which the communication takes place.

Affective Reasoning

Affective reasoning is concerned with an agent's *appraisal process*, where events are evaluated as to their emotional significance for the agent (Ortony *et al.* 1988, Elliott 1992). The significance is determined by so-called 'emotion-eliciting conditions', which comprise an event and three types of mental concepts.

- *Goals.* States of affairs that are (un)desirable, that the agent wants (does not want) to obtain.
- *Standards.* Beliefs about what ought (not) to be the case, events the agent considers as praiseworthy.
- *Attitudes.* Dispositions to like or dislike other agents or objects, what the agent considers appealing.

Those mental states are considered to be independent in the sense that an agent may be said to have a certain attitude without referring to its standards or goals. According to the emotion model of Ortony, Clore, and Collins (1988), also known as the OCC model, emotion types are just classes of eliciting conditions, each of which is labeled with an emotion word of phrase.[1] In total, twenty-two classes of eliciting conditions are identified. The OCC model can be represented by a set of rules and might thus be seen as a model of (emotion) *causation* that allows to reason about emotion.

One of the simplest emotions is the well-being emotion *joy* which has the following specification.

$$joy(L,F,\delta,S) \leftarrow wants(L,F,\delta_{Des(F)},S) \wedge holds(F,S)$$

In words, a (locutor-)agent L is in the emotional state of *joy* about fluent (i.e., state of affairs) F with intensity degree δ in situation S, if F is desirable for L in S with desirability degree $\delta_{Des(F)}$, and F holds in S. For all intensities of mental concepts related to reasoning about emotion, such as emotion intensities or goal desirability, we assume intensities $\delta \in \{1, ..., 5\}$. In the case of *joy*, we set $\delta = \delta_{Des(F)}$. In general, however, assigning appropriate intensities to emotions is a nontrivial task (Ortony, Clore, & Collins 1988), (Reilly 1996), (Gratch 2000). Consider the fortunes-of-others emotion *happy-for*, which is formalized as

$$\begin{aligned} &happy\text{-}for(L1,L2,F,\delta,S) \leftarrow \\ &\quad likes(L1,L2,\delta_{App(L1,L2)}) \wedge \\ &\quad joy(L2,F,\delta',S0) \wedge S0 < S \end{aligned}$$

where δ' is the presumed intensity of the *joy* emotion of the observed agent. For instance, if the observed agent $L2$ expresses happiness, which is communicated to the observing agent $L1$ in the form of a *com_act/6*

[1] Ortony *et al.* (1988) clearly distinguish between (emotion) types and (emotion) tokens, whereby the latter ones all share the specification of the corresponding type. E.g., the emotion type *joy* is associated with the tokens 'happy', 'cheerful', 'pleased', and so on. Having said this, we will subsequently often use 'emotion' rather than 'emotion types'.

representation, $L1$ has good reasons to believe that $L2$ is in the emotional state of *joy*. On the other hand, if the observing agent has beliefs about the observed agent's goals and their desirability, the agent can infer the emotional state of the other agent by using the very same emotion rules (see also Elliott and Ortony's (1992) *Concerns-of-Other* representations).

Following Reilly (1996), we employ logarithmic combination to compute the intensity of an emotion, i.e., for intensities δ and γ, the combined intensity is $\log_2(2^\delta + 2^\gamma)$. So, if an agent has evidence that another agent is very joyful ($\delta = 4$) and has a very positive attitude towards the other agent ($\delta_{App} = 5$), then the intensity of the agent's *happy-for* emotion would be 5 (computed values δ are rounded and set to 5 if $\delta > 5$). There are other ways to combine intensities (winner-takes-all, additive), but we found this choice the most natural. The specification of *happy-for* also assumes that the other agent was happy some time before the agent holds that belief. If needed, frame axioms are used to project facts to future states.

Prospect-based emotions such as *hope* or *disappointment* require calculating the probability of goal attainment, i.e., reasoning about plan states. Since our current model does not support this functionality, values have to be set in advance (but see Gratch (2000) for a thorough treatment of this issue). As another example, we shortly introduce the combined emotion *angry-at* (*reproach* and *distress*), which depends on the agent's standards.

$$\begin{aligned} &\textit{angry-at}(L1,L2,F,\delta,S) \leftarrow \\ &\quad \textit{holds}(\textit{did}(L2,A),S) \wedge \textit{causes}(A,F,S0) \wedge \\ &\quad \textit{wants}(L1, \neg F, \delta_{Des(\neg F)}, S) \wedge \textit{opposite}(F, \neg F) \wedge \\ &\quad \textit{blameworthy}(L1,A,\delta_{Acc(A)}) \wedge \textit{prec}(S0,S) \end{aligned}$$

Briefly, this means that an agent is angry with another agent if an undesirable fluent is caused by that agent's blameworthy action performed in the previous situation. δ_{Acc} refers to the degree to which the action is not acceptable for the agent.

Many models of emotion seem to suggest that once we have derived an agent's emotional state, all we have to do is to just let the agent express its emotion. However, it is far from clear how to express a *happy-for* emotion, and how to distinguish the expression of this emotion from the expression of a *joy* emotion or a *hope* emotion. Moreover, there might be no direct mapping, e.g., between the *angry-at* emotion and the expression of *anger*. At this point, the agent's personality comes into the play, as well as features of the social context. Personality and social setting will be the topic of the next section. Before that, we will briefly discuss the issue of emotion expression.

Emotions can be expressed through various channels, such as facial display, speech and body movement. The so-called 'basic emotions' approach (Ekman 1992) extracts those emotions that have distinctive (facial) expressions associated with them:[2] *fear*, *anger*, *sadness*, *happiness*, *disgust*, and *surprise*. Murray and Arnott (1995) describe the vocal effects on the basic emotions found in (Ekman 1992), e.g., if a speaker expresses *happiness*, his or her speech is typically faster, higher-pitched, and slightly louder. When running our human-agent conversation system, however, we found that vocal cues are rather ambiguous and therefore often rely on linguistic style to clearly express an agent's emotion.

Filter Programs

We have argued that goals, standards, and attitudes are the core mental concepts involved in an agent's appraisal of events, leading to a particular emotion. Besides 'internal' emotional states, we briefly discussed the agent's ('external') expression of its emotion. *Filter programs* are at the interface of the affective reasoner and the emotion expression module. They decide whether an emotion is expressed or suppressed, as well as the way and intensity in which an emotion is expressed. In our model, two factors determine emotion expression: the agent's personality and the agent's awareness of conventional practices that are applicable to the (social) situation.

Personality

Moffat (1997) suggests the following characterization:

> **Personality** is the name we give to those reaction tendencies that are consistent over situations and time. (Moffat 1997, p. 133)

Moffat argues that there is a close relationship between personality and emotions, although they seem very different at first sight. Emotions are short-lived and focused on particular events, whereas personality is stable and not focused. As a working hypothesis, Moffat assumes that the same cognitive structure underlies both emotion and personality. So, e.g., personality can be considered as a permanent and global emotion. In this paper, we prefer to keep mental concepts distinct, and use an agent's personality to bias an agent's emotion expression, given a certain emotional state. Thereby, we can guarantee that the agent's behavior is consistent (with its personality), which is of key importance to an agent's believability (Rousseau & Hayes-Roth 1998). In our current system, however, which only allows for rather short user-agent interactions, personality essentially collapses with *mood*, which is global (like personality) and rather short-lived (like emotions).

Our personality model is very simple, and considers just two dimensions, which seem crucial for social interaction.

- *Extroversion* refers to an agent's tendency to take action: sociable, active, talkative, optimistic.

[2] As there is only a limited number of comprehensive 'emotion words', we use *slanted* when referring to basic emotions rather than *italics* for emotional states.

- *Agreeableness* refers to an agent's disposition to be sympathetic: friendly, good-natured, forgiving.

In our model, we assume numerical quantification of dimensions, with a value from the set $\{-3, -2, -1, 1, 2, 3\}$. For instance, a value of 3 in the agreeableness dimension means that the agent is very friendly. Currently, the extroversion dimension is only considered in the conversational manager, where the initiative agenda (turn-taking) is handled.

Conventional Practices

A significant portion of human conversation takes place in a socio-organizational setting where participating agents have clearly defined *social roles*, such as sales person and customer, or instructor and student (Moulin 1998). *Conventional practices* are guidelines (or restrictions) about socially appropriate behavior in a particular social setting. We can distinguish two kinds of guiding restrictions:

- *Behavioral constraints* concern responsibilities, rights, duties, prohibitions, and possibilities associated with a social role.
- *Communicative conventions* function as a regulatory for the agent's choice of verbal expressions in a given social context.

Our interest is the choice of verbal and non-verbal behavior (emotion expression), depending on the agent's social role and personality.

Formally, in social or organizational groups, roles are ordered according to a *power scale*, which defines the social power of an agent's role over other roles (Moulin 1998). Power relations between agents L_i and L_j are represented as $P = p(L_i, L_j)$, where $P \in \{0, 1, 2, 3\}$. The value 0 means that agents have same rank. Otherwise, e.g., if the value of P is 1, then the rank of L_i is slightly higher than the rank of L_j. The *social network* is specified by the social roles and associated power relations. Following (Walker, Cahn, & Whittacker 1997), we also consider the *social distance* D between two agents, represented as $D = d(L_i, L_j)$, where $D \in \{0, 1, 2, 3\}$. If agents know each other very well, the social distance can be set to 0. Observe that the social distance between two agents can be high even if they have the same rank.

When agents interact, they do not only exchange information but also establish and maintain social relationships. Hence it is important that agents avoid introducing disharmony into a conversation (Moulin 1998) or threaten other agents' public face (Walker, Cahn, & Whittacker 1997). We assume that emotion expression (e.g., facial display or linguistic style) is determined by personal experience, background knowledge, and cultural norms (Walker, Cahn, & Whittacker 1997), as well as the 'organizational culture' (Moulin 1998). Consequently, human agents determine the values of the social variables 'social power' and 'social distance'. Based on the values of social power $P = p(L_j, L_i)$ and social distance $D = d(L_i, L_j)$, the agent L_i computes the *threat* θ from L_j of expressing a certain emotion by using the following simple equation

$$\theta = p(L_j, L_i) + d(L_i, L_j).$$

High values for θ lead to the agent's suppression of the expression of its emotional state, whereas low values allow the agent to show its emotional state. If $\theta = 0$, the agent considers itself as of same rank and high familiarity with the other agent. A zero value can also mean that the agent does not respect conventional practices. In (Prendinger & Ishizuka 2001c), we suggest the term *social role awareness* as a mental concept (or control state) that determines socially situated behavior.

Filter Rules

Basically, a social filter program consists of a set of rules that encode qualifying conditions for emotion expression. The program acts as a 'filter' between the agent's emotional state and its rendering in a social context, such as a conversation. As mentioned above, we consider the agent's personality and the agent's awareness of its social role as the most important emotion expression qualifying conditions. The most notable aspect of filter programs is that they introduce a level of *indirection* between an agent's emotion and the agent's expression of emotion. An agent may express anger although it is not angry but has a very unfriendly personality, or the agent may not express anger even if angry because it communicates with an agent that is socially more powerful.

In the following, we will give some examples of such filter rules. If the conversational partner has more social power or the social distance is large (i.e., θ is high), the expression of 'negative' emotions is typically suppressed, resulting in 'neutralized' emotion expression.

anger(L1,L2,ϵ,S) $\leftarrow$
social_threat(L2,L1,θ) $\wedge$
personality(L1,α) $\wedge$
angry-at(L1,L2,F,δ,S)

The first condition of the rule concerns social parameters, the second condition refers to the agent's agreeableness, and the third condition accounts for the output of the affective reasoner, the emotional state. As a first approximation, the intensity ϵ of emotion expression is computed as

$$\epsilon = \delta - (1 + \alpha + \theta)$$

As an example, consider the case where the agent is very angry (i.e., $\delta = 5$), rather unfriendly ($\alpha = -2$), and the social threat is maximal ($\theta = 6$). Here, $\epsilon = 0$, which means that the emotion is completely suppressed. On the other hand, if the agent does not respect social practices, i.e., $\theta = 0$, the agent's agreeableness dimension comes into effect, resulting in $\epsilon = 6\ (= 5 - (1 - 2))$. Since five is the maximal intensity level, greater values are cut off. If the agent is definitely angry ($\delta = 4$) but very friendly ($\alpha = 3$) and $\theta = 0$, then the *anger* intensity is zero, i.e., the agreeableness personality dimension 'consumes' the negative emotion.

Let us now discuss the effect of personality and social context on positive emotions. Observe that there are only two positive (universal) emotion expressions, *happiness* and (one interpretation of) *surprise*.

happiness(L1,L2,ϵ,S) $\leftarrow$
social_threat(L2,L1,θ) $\wedge$
personality(L1,α) $\wedge$
joy(L1,F,δ,S)

The intensity of positive emotions is computed as

$$\epsilon = \delta - (\theta - \alpha).$$

Consequently, an unfriendly personality or a high threat will diminish the expression of a positive emotion. Consider an agent that is very happy ($\delta = 5$) but unfriendly (e.g., $\alpha = -2$), communicating with a slightly distant conversant (i.e., $\theta = 1$). This agent will express *happiness* with intensity degree $\epsilon = 2$. The equations we currently use for computing the intensity of emotion expression are certainly not 'objective', although they seem to bear some plausibility. People who use our system (as autonomous agent designers) are free to modulate the parameters.

As demonstrated in (Prendinger & Ishizuka 2001a; 2001c), users interact with our system by uttering one of a set of pre-defined sentences with varying linguistic style.

- I would like to drink a cup of coffee. (*polite*)
- I will have a cup of coffee. (*neutral*)
- Bring me a coffee, right away. (*rude*)

An interesting phenomenon of human-human communication is the *reciprocal feedback loop* where, e.g., one agent's use of polite linguistic style results in another agent adapting its linguistic style. We support a limited form of feedback, depending on the user's (or agent's) linguistic style, by adding or subtracting one intensity 'unit'. Hence, if the agent would give a cheerful answer with intensity degree $\epsilon = 4$, it will respond with degree 5 if asked politely, and with degree 3 if asked in a rude way. A neutral question does not change the emotion expression intensity.

Summary of Empirical Study

We conducted an experiment that investigates the impact of animated agents featuring social role awareness (see Prendinger and Ishizuka 2001b for more details). In the experiment, participants would play the role of a customer in a virtual coffee shop and interact with an animated agent portraying a waiter. The waiter agent interacts with a manager agent and another customer agent that turns out to be an old acquaintance of the waiter. All agents are characters available for the Microsoft Agent package ("James", "Genie", and "Al"). The participants (16 in total) promoted the conversation by simply clicking a radio button next to a predefined conversational contribution. In order not to distract participants from the agents' reactions, speech recognition was disabled in the interaction.

We prepared two different versions of the system that were identical except for the following features:

- In the *Unfriendly Waiter* version ($C1$) the waiter agent James responded to the user in a rude way, but changed to friendly behavior when interacting with his manager and the other customer (an old friend). E.g., if the user orders alcohol, James would say "Get out of here. This is a coffee-shop!". On the other hand, James would warmly welcome the other customer (Al), after he recognizes him as an old friend.
- In the *Friendly Waiter* version ($C2$), the waiter agent displayed polite behavior to the user but disobeyed the manager's order and turned down his old friend.

As to the naturalness or believability of James' behavior, we hypothesized the following outcome of the experiment. In the *Unfriendly Waiter* version ($C1$), subjects would rate James' behavior as unnatural towards themselves (as customers) but natural towards the other agents (manager, friend). Moreover, they would think that in general, James has an unfriendly (disagreeable) personality. In the *Friendly Waiter* version ($C2$), on the other hand, subjects would consider James' behavior natural towards themselves but inappropriate towards the other agents, and would find James' personality friendly.

It could be shown that subjects considered James' behavior towards users significantly more natural in the $C2$ version than in the $C1$ version. Moreover, subjects found James' behavior significantly more agreeable in the $C2$ version than in the $C1$ version. This result is interesting since we set up the system in a way that James displays (un)friendly behavior about half of the total interaction time. Hence subjects seemed to distinguish between behavior that is influenced by the agent's personality and behavior that is motivated by the particular social role the agent takes in a particular situation. However, contrary to our expectation, subjects considered James' behavior towards other agents slightly more natural in the $C2$ version than in the $C1$ version, although James ignored behavioral practices towards other agents in the $C2$ version. A possible reason is that James' interaction with the other agents was too short to clearly estimate the appropriateness of James' responses.

Discussion and Conclusion

In this paper, we propose to include reasoning about personality and social context to mental models of animated agents, which complements the affective reasoning component. All of the mental concepts (control states) involved in the reasoning process have associated intensities. Our initial experience with a web-based system implementing those features indicate enhanced believability of animated characters, at least for language conversation training tasks (Prendinger & Ishizuka 2001c; 2001b). However, our system can also be used as a general-purpose platform for experimenting with behavioral patterns of animated agents, which

is important for evaluating users' reactions to different styles of agent behavior.

The prime focus of the proposed filter program is to support high social accuracy of agents' responses, e.g., by considering parameters such as social power and distance. For user-adapted interaction, the animated agents can be customized in various ways by modifying the parameters underlying their behavior. However, we currently have no principled way to introduce *social feedback* mechanisms to our model, i.e., values for social power and distance are kept fixed. On the other hand, when humans establish and maintain social relationships, e.g., the value of their social distance changes. This phenomenon is not reflected in agents' mental models. Agents have no memory about the social implications of their interactions with a human or other agents. We can imagine that if a user repeatedly interacts with an agent in a friendly way, the agent should not only remember the user but also show more familiarity, by choosing appropriate verbal and non-verbal communication acts. An exciting direction for future research might be to incorporate a reinforcement-based algorithm for enhanced agent-user adaption (Breazeal & Velásquez 1998).

Acknowledgments

We would like to thank the reviewers for their helpful and valuable suggestions. This research is supported by a JSPS Research Grant (1999–2003) for the Future Program.

References

Allen, S. 1999. Control states and motivated agency. In André, E., ed., *Proceedings of the i3 Spring Days'99 Workshop on Behavior Planning for Life-like Characters and Avatars*, 43–61.

André, E.; Rist, T.; van Mulken, S.; Klesen, M.; and Baldes, S. 2000. The automated design of believable dialogue for animated presentation teams. In Cassell, J.; Prevost, S.; Sullivan, J.; and Churchill, E., eds., *Embodied Conversational Agents*. The MIT Press. 220–255.

Breazeal, C., and Velásquez, J. 1998. Toward teaching a robot "infant" using emotive communication acts. In *Proceedings of Workshop on Socially Situated Intelligence*.

Cassell, J., and Thórisson, K. R. 1999. The power of a nod and a glance: Envelope vs. emotional feedback in animated conversational agents. *Applied Artificial Intelligence* 13(4–5):519–538.

de Rosis, F.; de Carolis, B.; and Pizzulito, S. 1999. Software documentation with animated agents. In *Proceedings 5th ERCIM Workshop on User Interfaces For All*.

Ekman, P. 1992. An argument for basic emotions. *Cognition and Emotion* 6(3–4):169–200.

Elliott, C., and Ortony, A. 1992. Point of view: Modeling the emotions of others. In *Proceedings 14th Annual Conference of the Cognitive Science Society*, 809–814.

Elliott, C. 1992. *The Affective Reasoner. A process model of emotions in a multi-agent system*. Ph.D. Dissertation, Northwestern University. Institute for the Learning Sciences.

Gratch, J. 2000. Émile: Marshalling passions in training and education. In *Proceedings 4th International Conference on Autonomous Agents (Agents'2000)*, 325–332.

Ishizuka, M.; Tsutsui, T.; Saeyor, S.; Dohi, H.; Zong, Y.; and Prendinger, H. 2000. MPML: A multimodal presentation markup language with character control functions. In *Proceedings Agents'2000 Workshop on Achieving Human-like Behavior in Interactive Animated Agents*, 50–54.

Johnson, W. L.; Rickel, J. W.; and Lester, J. C. 2000. Animated pedagogical agents: Face-to-face interaction in interactive learning environments. *International Journal of Artificial Intelligence in Education*.

Microsoft. 1998. *Developing for Microsoft Agent*. Microsoft Press.

Moffat, D. 1997. Personality parameters and programs. In Trappl, R., and Petta, P., eds., *Creating Personalities for Synthetic Actors*. Springer. 120–165.

Moulin, B., and Rousseau, D. 2000. An approach for modeling and simulating conversations. In Vanderveken, D., and Kubo, S., eds., *Essays in Speech Act Theory*. John Benjamins.

Moulin, B. 1998. The social dimension of interactions in multiagent systems. In Wobcke, W.; Pagnucco, M.; and Zhang, C., eds., *Agents and Multi-Agent Systems, Formalisms, Methodologies and Applications*. Springer Verlag. 109–122.

Murray, I. R., and Arnott, J. L. 1995. Implementation and testing of a system for producing emotion-by-rule in synthetic speech. *Speech Communication* 16:369–390.

Ortony, A.; Clore, G. L.; and Collins, A. 1988. *The Cognitive Structure of Emotions*. Cambridge University Press.

Paiva, A.; Machado, I.; and Prada, R. 2001. Heroes, villains, magicians,...: Dramatis personae in a virtual story creation environment. In *Proceedings International Conference on Intelligent User Interfaces (IUI-2001)*, 129–136.

Prendinger, H., and Ishizuka, M. 2001a. Carrying the role-playing metaphor to interactive learning environments. In *International Conference on Intelligent User Interfaces (IUI-2001)*.

Prendinger, H., and Ishizuka, M. 2001b. Let's talk! Socially intelligent agents for language conversation training. *IEEE SMC Transactions, Special Issue on "Socially Intelligent Agents"*. To appear.

Prendinger, H., and Ishizuka, M. 2001c. Social role awareness in animated agents. In *Proceedings 5th International Conference on Autonomous Agents (Agents-01)*. To appear.

Reilly, W. S. N. 1996. *Believable Social and Emotional Agents*. Ph.D. Dissertation, Carnegie Mellon University. CMU-CS-96-138.

Rousseau, D., and Hayes-Roth, B. 1998. A social-psychological model for synthetic actors. In *Proceedings 2nd International Conference on Autonomous Agents (Agents'98)*, 165–172.

Walker, M. A.; Cahn, J. E.; and Whittacker, S. J. 1997. Improvising linguistic style: Social and affective bases for agent personality. In *Proceedings First International Conference on Autonomous Agents (Agents'97)*.