

Appraisal and Filter Programs for Affective Communication

Helmut Prendinger and Mitsuru Ishizuka

Department of Information and Communication Engineering
Graduate School of Information Science and Technology, University of Tokyo
7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
{helmut,ishizuka}@miv.t.u-tokyo.ac.jp

Abstract

This paper introduces a model of user-agent and inter-agent interaction that supports basic features of affective communication. As essential requirements for animated agents' capability to engage in social interaction, we motivate reasoning about emotion and emotion expression, personality, and social role awareness. A (rather standard) appraisal program is employed to derive the agent's emotional state. The novel aspect of our approach is the introduction of a filter program that qualifies the agent's expression of its emotional state by its personality and more importantly, by the social setting in which the conversation takes place. This allows an agent to suppress an emotion, if the expression of the emotion would defeat a higher-order goal. We also discuss rudimentary mechanisms of social feedback.

Introduction

Our concept of affective communication is motivated by the influential paradigm of *affective computing*, "computing that relates to, arises from, or deliberately influences emotions" (Picard 1997). In line with this work, we assume that emotions are indispensable for effective communication, and thus promote the view that emotions should be integrated to models of human-computer interaction. Specifically, we envision that humans interact with animated agents, e.g., cartoon-style characters that may behave in believable and socially appropriate ways.

Recent years have seen a growing body of literature that aims to integrate emotions to architectures of autonomous agents. In order to position our own work, we will broadly categorize emotion research for autonomous agents into *cognitive* emotion approaches and *non-cognitive* (or innate) emotion approaches. Research in the first category mainly derives from Ortony, Clore, and Collins' (1988) seminal work in cognitive psychology, where emotions are seen as valenced (i.e., positive or negative) reactions to events such as other agents' actions or objects, relative to the agent's goals, standards, and attitudes (Elliott 1992; Reilly 1996; Gratch 2000). On the other hand, non-cognitive emotion approaches are mostly inspired by theories in neu-

rosience and ethology, which emphasize 'low-level' influences to emotion generation besides cognitive reasoning. For instance, Velásquez (1998) considers four elicitors in his emotion model: neural, sensimotor processes, motivations, and cognition (see also Breazeal and Velásquez (1998)).

Given this (rather ad hoc) classification, our approach clearly falls into the first category that approaches emotion from the cognitive point of view. Our main interest is the affective dimension in communication, which can be best modelled by explicitly representing mental concepts. In particular, our goal is to develop a general framework for affective and social communication that covers various forms of human-agent and inter-agent interactions. Consequently, we will propose 'higher-level' components of agents' mental models that enable them to process emotions and show affective behavior. A salient feature of our model is the following distinction:

- *Emotional states* are the result of reasoning about events, an agent's goals, standards, and attitudes.
- *Emotion expression* is the result of reasoning about the agent's emotional state, qualified by the agent's personality (or mood) and the social context.

A consequence of this distinction is a level of indirection between emotional state and emotion expression, which is mandatory when agents act in social settings where conventional practices apply (Moulin & Rousseau 2000). So, for instance, an angry agent might not show its anger, if the agent interacts with another agent that has more social power or to which the social distance is large. Corresponding to the before mentioned distinction, two mechanisms are employed. An *appraisal program* evaluates an event as to its emotional significance for the agent, whereas a *filter program* qualifies the agent's emotion by its personality and the agent's awareness of the social threat from the other agent. For illustration, consider that someone crashes your computer. Depending on your attitude towards computers, you might be in a happy or angry emotional state. However, whether you will show your emotion will largely depend on your agreeableness and the social position of the aggressor. We claim that mechanisms of expression

and suppression of emotional states are of key importance for socially intelligent behavior. Given a high-order maintenance goal, e.g., “keep your job”, an agent might suppress its ‘angry’ state, if it happens to be his or her boss who crashes his or her computer. The phenomenon of emotion suppression also gained much interest in the psychology literature recently, where it is called *emotion regulation* (Gross 1998).

In this paper, we will describe the influence of mental concepts to emotional state and emotion expression. Important issues include the integration of the intensities of the various mental 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 we adhere to the “conversation training as role-playing in interactive games” metaphor as an enjoyable learning environment. The programmable interface of the Microsoft Agent package is used to run our interactive role-playing scenarios. The package comes ready with a speech recognizer and text-to-speech engine that allow client-side execution in a web browser. We currently use off-the-shelf 2D cartoon-style characters, but have 3D agents under development to overcome restrictions on the embodied behavior of the 2D agents. Our system is discussed in a related paper (Prendinger & Ishizuka 2001a).

The rest of the paper is organized as follows. In the next section, we sketch our communication model and interaction protocol. After that, we describe an affective reasoner for reasoning about emotion, and report on an influential theory of emotion expression. The following section introduces a so-called ‘filter program’ that functions as a ‘filter’ between 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. Next, we outline two simple social feedback mechanisms. The final section concludes the paper.

Modelling Conversations among Autonomous Agents

Conversations can be characterized by the presence of multiple (locutor-)agents that communicate through various channels, such as verbal utterances, gestures, body movement, and facial display. Following Moulin and Rousseau (2000), we can 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 this paper, we will focus on the conceptual and the social level.

At the conceptual level, information is exchanged between agents as simplified symbolic representations of the utterance together with stylistic and affective mark-

ers. Consider an agent called ‘Al’ ordering beer from an agent called ‘James’, by saying “Bring me a beer, right away”. As a basic interaction protocol between agents, we propose communicative acts of the form

com_act(al,james,order_beer,rude,anger,s0)

where ‘al’ is the speaker, ‘james’ is the hearer, ‘order_beer’ is the conveyed information, ‘rude’ is a qualitative evaluation of the linguistic style of the utterance, ‘anger’ refers to Al’s emotion expression, and *s0* denotes the situation in which the utterance takes place. In animated agents, the expressed emotion can be generated, e.g., by acoustic signals (Murray & Arnott 1995) or facial display (Ekman 1992). Concerning users’ emotion expression, ‘neutral’ is set as the default value, as an emotion (expression) recognition module is not part of our system (but see Picard (1997)).

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, social role awareness). Similarly, 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., attitude), it is a good indicator that mood is a control state. The integration and interaction of those concepts allows for a broad variety of believable agent behaviors that might be conceived as intelligent.

Appraisal Program

Affective reasoning is concerned with an agent’s *appraisal process*, where events are evaluated as to their emotional significance for the agent (Ortony, Clore, & Collins 1988; Elliott 1992). The significance is determined by so-called ‘emotion-eliciting conditions’ (EECs), which comprise an event and three types of mental concepts: (i) *goals*, i.e., states of affairs that are (un)desirable, what the agent wants (does not want) to obtain; (ii) *standards*, i.e., the agent’s beliefs about what ought (not) to be the case, events the agent considers as praiseworthy; and (iii) *attitudes*, i.e., the agent’s dispositions to like or dislike other agents or objects, what the agent considers appealing.

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 labelled with an emotion word or phrase.¹ In

¹Ortony *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’,

total, twenty-two classes of eliciting conditions are identified. One of the simplest emotions is the well-being emotion *joy* which has the following specification.

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

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 . We assume intensities $\delta \in \{0, \dots, 5\}$, whereby, e.g., the zero value for an emotion type means that the threshold for activating the associated emotion has not been crossed. 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

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

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* 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 (default) 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).²

Following Reilly (1996), we employ logarithmic combination to compute the intensity of an emotion, i.e., for intensities δ_i , the combined intensity δ is $\log_2(\sum_i 2^{\delta_i})$. 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, situation calculus 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

¹'cheerful', 'pleased', and so on. Having said this, we will subsequently often use 'emotion' rather than 'emotion types'.

²Note that we assume that all EECs have the status of 'propositional attitudes', i.e., relations between an agent and a mental concept (e.g., 'beliefs'). Those concepts comprise such seemingly diverse entities as what the agents wants or the agent's assessment of another agent's emotional state.

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.

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

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:³ *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 Program

Following the OCC model, 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 emotional state. Besides 'internal' emotional states, we briefly discussed the agent's ('external') expression of its emotion. A *filter program* is at the interface of the affective reasoner and the emotion expression module. It decides whether an emotion is expressed or suppressed, as well as the way and intensity in which an emotion is expressed. The filter program closely resembles the regulation rules of De Carolis *et al.* (2001). In our model, two factors determine emotion expression: the agent's personality and the agent's

³As there is only a limited number of comprehensive 'emotion words', we use *slanted* when referring to basic emotions rather than *italics* for emotional states.

awareness of conventional practices that are applicable to the (social) situation.

Personality

Moffat (1997) suggests to characterize *personality* as “the name we give to those reaction tendencies that are consistent over situations and time” (p. 133). Whereas emotions are short-lived and focused on particular events, personality is stable and not focused. Personality guarantees that an agent will behave consistently which is considered of key importance to the agent’s believability (Rousseau & Hayes-Roth 1998).

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. *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 $\{-5, \dots, 5\}$. E.g., a value of 5 in the agreeableness (extroversion) dimension means that the agent is very friendly (very introvert). In general, we consider personality as a ‘regulation’ condition that estimates to what extent the agent is able (or willing) to control its behavior. In general, *regulation* ρ is computed as

$$\rho = \frac{\sum_i \rho_i}{N}$$

where the denominator N scales the result according to the number of considered control parameters. Currently, we only consider the agent’s agreeableness ρ_A and extroversion ρ_E . Note that the equation simply captures the intuition that one dimension of personality may defeat another one, e.g., an agent that is unfriendly but introvert manifests higher regulation than an unfriendly and extrovert agent.

Let us briefly show how our personality concept compares to Elliott’s (1992) distinction between *interpretive* personality (what concerns the agent) and *manifestative* personality (how the agent reacts). Elliott considers interpretive personality as an umbrella term for the EECs (discussed above), while we take it as a separate concept where, e.g., what is desirable for the agent is ‘caused’ by its personality. On the other hand, Elliott’s concept of manifestative personality is very similar to ours but more fine-grained as he provides a specific set of behaviors for each agent.

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 (see also Ekman and Friesen’s (1969) notion of ‘display rules’). 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 $\theta_P = p(L_i, L_j)$, where $\theta_P \in \{0, \dots, 5\}$. The value 0 means that agents have same rank. Otherwise, e.g., if the value of θ_P is 2, 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* θ_D between two agents, represented as $\theta_D = d(L_i, L_j)$, where $\theta_D \in \{0, \dots, 5\}$. 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 $\theta_P = p(L_j, L_i)$ and social distance $\theta_D = d(L_i, L_j)$, the agent L_i computes the *threat* θ from L_j of expressing a certain emotion by using the following simple equation

$$\theta = \sum_i 2^{\theta_i}$$

High values for θ typically 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 ignores conventional practices. Observe that our list of ‘threats’ is by no means exhaustive. Besides power and distance, De Carolis *et al.* (2001) also consider, e.g., the interlocutor’s personality and cognitive capacity, and the type of social situation (private or public). Those parameters can be easily added to our model. In Prendinger and Ishizuka (2001b), 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.

In the following, we will give some examples of such filter rules. If the social threat is high, the expression of ‘negative’ emotions is typically suppressed, resulting in ‘neutralized’ emotion expression (our distinction into ‘positive’ and ‘negative’ emotions is based on Reilly’s (1996) categorization).

$$\begin{aligned} \text{anger}(L1, L2, F, \epsilon, S) \leftarrow \\ \text{social_threat}(L2, L1, \theta) \wedge \\ \text{regulation}(L1, \rho) \wedge \\ \text{angry-at}(L1, L2, F, \delta, S) \end{aligned}$$

The first condition of the rule concerns social parameters, the second condition refers to the agent’s regulation capacity (e.g., its 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 $\epsilon = \delta - (\theta + \rho)$ where $\epsilon \in \{0, \dots, 5\}$. As an example, consider the case where the agent is very angry (i.e., $\delta = 4$), rather unfriendly ($\rho_A = -2$), and the social threat is high ($\theta = 6$). 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 $\epsilon = 6 (= 4 - (-2))$. Since five is the maximal intensity level, values above this threshold are cut off. If the agent is definitely angry ($\delta = 4$) but very friendly ($\rho_A = 4$) and $\theta = 0$, then the *anger* intensity is zero, i.e., the agent’s agreeableness ‘consumes’ the negative emotion. Observe that the suppression of a negative emotional states such as *angry-at* can be called ‘intelligent’ if it allows the agent to uphold certain goals which would otherwise be threatened by the direct expression of the negative emotion.

In the following, we briefly discuss the effect of personality and social context on the expression of positive emotions, such as *joy*.

$$\begin{aligned} \text{happiness}(L1, L2, F, \epsilon, S) \leftarrow \\ \text{social_threat}(L2, L1, \theta) \wedge \\ \text{regulation}(L1, \rho) \wedge \\ \text{joy}(L1, F, \delta, S) \end{aligned}$$

The intensity of positive emotions is computed as $\epsilon = \delta - (\theta - \rho)$. Consequently, the agent’s 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., $\rho_A = -2$), 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.

Social Dynamics

Although filter programs endow agents with awareness of the social context they are situated in, they do not provide them with mechanisms to *change* social parameters as a result of social interaction. However, when (human) agents interact, they establish and maintain social relationships, hence, e.g., the value of social distance changes, an issue that attracted strong interest in the biology-inspired field of Socially Intelligent Agents (see Dautenhahn (2000), Cañamero (2000)). In the following, we will discuss an instance of short-term as well as long-term social feedback.

Reciprocal Feedback

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. For instance, consider the following utterances with varying linguistic style (polite, neutral, rude):

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

Our system supports a limited form of reciprocal feedback, whereby depending on the user’s (or agent’s) linguistic style, ‘intensity units’ are added or subtracted to (from) the agreeableness degree. Hence, if the agent would give a cheerful answer with intensity degree $\epsilon = 3$, it might respond with degree 5 if asked politely, and with degree 1 if asked in a rude way (given appropriate intensity values for the remaining control states). A neutral question does not change the emotion expression intensity.

Social Feedback

As opposed to the rather momentary influence of reciprocal feedback, *social feedback* refers to changes of agents’ relationships over more extended periods of time. For instance, when agents interact frequently, the value of social distance will eventually decrease, and also, power relations can change. Similarly, the appeal-iness degree associated with the agent’s attitude may undergo changes when the agent repeatedly faces positive (negative) feedback from another agent. A promising direction for future research might be to incorporate a reinforcement-based algorithm for enhanced agent-user adaption (Breazeal & Velásquez 1998).

In fact, the very basis of dramatic action is that values associated with mental concepts turn (change). We started to investigate the mechanisms behind social dynamics resulting from interactions, which will allow for a richer interaction experience than one solely based on

agents' believable reactions. Currently, intensities of control states are manipulated in a rather ad-hoc fashion, by summing up positive (negative) feedback values and updating the overall intensity of mental concepts.

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 2001a). 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 an agent's affective responses.

The prime focus of the proposed filter program is to support high social accuracy (or appropriateness) of agents' responses, e.g., by considering parameters such as social power and distance. The socially qualified response is motivated, e.g., as the suppression of a typically negative emotion in order to uphold a higher-level goal. By considering more sophisticated mechanisms of social interaction, we hope to gain a better understanding of the complexity of social interactions.

Acknowledgments

We would like to thank the reviewers for their valuable comments and helpful suggestions. The first author also greatly benefited from discussions with Fiorella de Rosis. 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.
- Breazeal, C., and Velásquez, J. 1998. Toward teaching a robot "infant" using emotive communication acts. In *Proceedings of Workshop on Socially Situated Intelligence*.
- Cañamero, D. 2000. Designing emotions for activity selection. Technical Report DAIMI PB545, Department of Computer Science, University of Aarhus, Denmark.
- Dautenhahn, K. 2000. Socially intelligent agents and the primate social brain. towards a science of social minds. In Dautenhahn, K., ed., *Socially Intelligent Agents: The Human in the Loop (Papers from the 2000 AAAI Fall Symposium)*, 35–51.
- de Carolis, B.; Pelachaud, C.; Poggi, I.; and de Rosis, F. 2001. Behavior planning for a reflexive agent. In *Proceedings 17th International Conference on Artificial Intelligence (IJCAI-01)*.
- Ekman, P., and Friesen, W. V. 1969. The repertoire of nonverbal behavior: Categories, origins, usage, and coding. *Semiotica* 1:49–98.
- 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.
- Gross, J. J. 1998. The emerging field of emotion regulation: An integrative review. *Review of General Psychology* 2(3):271–299.
- Moffat, D. 1997. Personality parameters and programs. In Trapp, 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.
- Picard, R. W. 1997. *Affective Computing*. The MIT Press.
- Prendinger, H., and Ishizuka, M. 2001a. Let's talk! Socially intelligent agents for language conversation training. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans, Special Issue on "Socially Intelligent Agents"* 31(5). To appear.
- Prendinger, H., and Ishizuka, M. 2001b. Social role awareness in animated agents. In *Proceedings 5th International Conference on Autonomous Agents (Agents-01)*, 270–277.
- 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.
- Velásquez, J. D. 1998. When robots weep: Emotional memories and decision-making. In *Proceedings 15th National Conference on Artificial Intelligence (AAAI-98)*, 70–75.
- Walker, M. A.; Cahn, J. E.; and Whittaker, S. J. 1997. Improvising linguistic style: Social and affective bases for agent personality. In *Proceedings First International Conference on Autonomous Agents (Agents'97)*.