Broadening the Scope of Affect Detection Research

Rainer Reisenzein

Abstract—I propose to broaden the scope of affect detection research in three related directions. First, the task definition should be broadened from the detection of affects to the inference of mental states. Second, the detection process should be reconceptualized as an inference to the best explanation that makes essential use of a theory of mind. Third, additional data sources should be utilized to infer emotions: information about the situation (eliciting events), information about emotion-related goal-directed actions, and self-reports about emotions.

Index Terms—Human-computer interaction, affective computing, affect detection, theory of mind, inference to the best explanation, emotion self-reports.

1 Introduction

With their interdisciplinary review of affect detection (AD), Calvo and D'Mello (this issue) have done a valuable service to both emotion scientists and affective computing researchers (see also, [38]). This comment is written from the perspective of an emotion psychologist. I first summarize the potential benefits of the recent AD research for emotion psychology, and then make three suggestions for how this research could be developed further.

2 How Affect Detection Research Can Benefit Emotion Psychology

The aim of the studies on affect detection by affective computing researchers is not to answer theoretical or empirical questions of emotion psychology; rather, it is the practical goal of increasing the quality of human-computer interaction by endowing computers (or, better, artificial agents) with the capacity to recognize, and react appropriately to, the emotions of humans (see [19]). Still, this research can provide several benefits to emotion psychology. First, it yields empirical data relevant for the questions of emotion psychology, such as data on which emotions are expressed in which behaviors under which conditions. Second, affective computing researchers have recently begun to extend the earlier psychological AD studies in several directions. The most important of these extensions are: the consideration of "nonbasic" emotions (e.g., frustration, interest, boredom) that have been largely neglected in psychological AD research, the study of spontaneous emotional expressions in everyday life situations and, related to this, the construction of new affect expression databases, the increasing consideration of multimodal affect detection, and the application of advanced computational methods for affect inference. Some of the algorithms proposed for this purpose could be of interest to psychologists as possible models of how humans infer emotions. But even when this is not the case, reliable automatic AD systems would be highly useful to emotion psychologists. For example, facial expression researchers would welcome an automated facial action coding system that relieves them of the chore of hand-coding videos (e.g., [22]).

• *The author is with the Institute of Psychology, University of Greifswald, Franz-Mehring-Strasse 47, 17487 Greifswald, Germany.*
E-mail: rainer.reisenzein@uni-greifswald.de.

Manuscript received 11 Aug. 2010; accepted 14 Aug. 2010; published online 23 Aug. 2010.
Recommended for acceptance by J. Gratch.
For information on obtaining reprints of this article, please send e-mail to: tac@computer.org, and reference IEEECS Log Number TAFFC-2010-08-0061.
Digital Object Identifier no. 10.1109/T-AFFC.2010.9.

The recent trends in AD research documented by Calvo and D'Mello are all on the right track, and I applaud them. However, I would like to persuade affective computing researchers to broaden the scope of their inquiry still further, in three closely related directions. First, I propose to *broaden the task definition* from the detection of affects to the inference of mental states. Second, I propose to *reconceptualize the detection process* as an inference to the best explanation that makes essential use of a theory of mind (and possibly simulation). Third, I propose to *utilize additional data sources* to infer affect: information about the situation (emotion-eliciting events), information about emotion-related goal-directed actions, and self-reports about emotions (and emotion-related mental states). These proposals converge on the suggestion to make the process of affect detection by artificial agents more human-like and, as a consequence, its outcomes more useful, both to humans and to the artificial agents themselves.

3 From Affect Detection to Mental State Inference

Affect detection is but a special case of the much more general problem of mindreading—the inference of others' mental states. Other than the historically contingent definitions of the field and tasks of affective computing, I can see no reason to restrict the detection problem to the inference of affective states. On the contrary, two arguments speak against such a restriction.

First, knowledge about nonaffective mental states, in particular cognitive states (beliefs) and motivational states (desires, intentions), is at least as important for enhancing the quality of human-computer interaction as is knowledge about emotions (e.g., [1]). Second, even if the focus is on the inference of emotions, cognitive and motivational states should be considered because emotions are closely connected to them (e.g., [18], [23]). As a consequence, if one has a theory of the relations of emotions to beliefs and desires, one can use this theory to infer emotions from knowledge about these other mental states (for an affective computing illustration of this idea, see [2]). Such knowledge, at least in the form of plausible conjectures, is indeed often available. For example, knowing or assuming 1) that you wanted to gain a point in a game (which I could simply infer from my general knowledge that most game-players want to gain points) and believed that you would gain it, and 2) that you have come to believe that the desired event did not materialize (which I could simply infer from seeing that you failed to gain the point), I can infer that you feel disappointed.

If affective computing researchers decide to broaden the task definition from affect detection (AD) to mental state detection (MSD), they may wish to consult the broader literature on mindreading that has accumulated in several fields of psychology. In particular, in social psychology, the study of common-sense psychology—which includes the problem of mindreading—has been a productive research area for more than 50 years ([10], for reviews see [16], [27]). In developmental psychology, the nature and ontogenetic development of the capacity to ascribe mental states to others has become the topic of an intense interdisciplinary discussion between psychologists and philosophers (see [4] for a recent review). This research has apparently been largely neglected so far in the field of affect detection, as has been related research by computational scientists, for example on intention recognition (e.g., [1]) and theory-of-mind-based reasoning (e.g., [21]).

4 Reconceptualizing Mental State Detection as a Theory-of-Mind Based Inference to the Best Explanation

What kind of process is MSD? The most plausible answer to this question, in my view, is that MSD is a process of *abduction*, or *inference to the best explanation*, in which that one of several possible mental states is attributed that fits best with the totality of the available evidence (see [25] for an application of this idea to the inference of surprise). Several computational models of abduction have been proposed (see, e.g., [13], [34]). Independent of which of these models is to be preferred, an important insight of Heider's [10] analysis of common-sense psychology is that MSD makes essential use of a theory of mind that specifies the typical connections between different mental states, situations (events), and behaviors. This insight connects up with the second reason I gave for my earlier suggestion to broaden the task definition of AD to MSD—that doing so allows the theory-based inference of emotions from other mental states. Now, however, the emphasis is on another benefit of theory-based inference: By placing constraints on the possible solutions to an MSD problem, a theory of mind helps to solve the problem that mental state inferences are usually underdetermined by the available data.

The proposed reconceptualization of MSD as a case of theory-guided inference to the best explanation entails that artificial agents need to be endowed with a (at least rudimentary) theory of the human mind, specifically the emotion system, to efficiently infer emotions. In addition, artificial agents also need an emotion theory for another reason, namely, to be able to react appropriately to the affect they detect in humans (a second main task of affective computing; see Calvo and D'Mello, this issue and [19]). Lacking an emotion theory—a theory that relates emotions not only to behaviors (e.g., facial expressions), but also to other mental states—the agent's emotion concepts (e.g., "fear," "anger") are extremely rudimentary: About all they allow the agent to do is to classify certain behaviors into a number of labeled categories. Not knowing the (likely) causes and consequences of a detected emotion, the agent is unable to systematically influence this emotion and to predict its effects on cognition and action.

In sum, without a theory of mind (that includes a theory of emotion), the capabilities of artificial agents to both detect emotions and to respond to them are extremely limited. In fact, unless an artificial agent's emotion concepts are embedded into a theory of mind that imbues them with (additional) meaning, one cannot even properly say that the agent detects emotions. The inference of emotion deserves to be called an inference of emotion *by the agent* only if the agent possesses a reasonable approximation to the emotion concepts of humans (e.g., "afraid," "happy"), and human emotion concepts acquire their meaning as terms in an implicit theory of mind.

The proposed reconceptualization of the process of affect detection as a case of theory-of-mind-based inference to the best explanation, as well as the related proposal to infer emotions from other mental states connected to them, imply that affective computing researchers need to pay attention to psychological theories of emotion (or, perhaps, to the emotion theory implicit in common-sense psychology), even if they are not interested in emotion psychology per se but only want to improve the affect detection and responding abilities of artificial agents.

The psychological and philosophical literature on MSD also suggests the existence, in humans, of an alternative or supplement to the theory-guided inference of emotions: Instead of using a theory of mind, one can infer the emotions of others by engaging in a mental simulation that makes use of one's own emotion system (see, e.g., [9]). For example, to find out how you might have felt at a certain point in a game, I can imagine myself being in your situation and observe how I feel. Humans probably make use of both processes: theory-guided inference and simulation (see also, [33]). Analogously, future artificial agents might infer or predict the emotions of humans by using their own artificial emotion mechanisms to simulate how they would feel in particular situations. Doing so presupposes three things: that the agents have emotion mechanisms, that these mechanisms are sufficiently human-like, and that the agents can use them "offline," i.e., for the purpose of simulation. If this suggestion is followed, the task of affect detection would become intimately connected to a third main task of affective computing ([19]), that of endowing artificial agents with the capacity to have emotions themselves.

5 Including Information about Eliciting Events, Emotional Actions, and Verbal Communications about Emotions

My third suggestion is closely connected to the first two: Affective computing researchers should broaden the range of the data sources from which affect is inferred. As Calvo and D'Mello note, so far the focus has been on the involuntary behavioral manifestations of emotions (facial expressions, prosodic features of speech, peripheral physiological changes). What has been largely neglected so far (noteworthy exceptions, concerning situational information, are [2], [12]) are three other sources of information about emotion that play an important role in both everyday and scientific emotion inference (see [25], for an example): 1) information about the situational context, in particular about emotion-eliciting events, 2) information about emotion-related goal-directed actions (e.g., aggression in the case of anger, helping in the case of pity, apologizing in the case of guilt; see [37]), and 3) verbal communications (self-reports) about emotions and emotion-related mental states. To see why these additional sources of information about affect are not just useful but necessary for successful AD, I digress to elaborate on the limitations of nonverbal emotion indicators to which Calvo and D'Mello (this issue) briefly refer in their discussion.

6 Limitations of Affect Detection from Nonverbal Expressions of Emotion

Affect inference from nonverbal manifestations of emotions is based on the premise that emotional states manifest themselves reliably in nonverbal behaviors (e.g., facial expressions, physiological changes). However, even the proponents of "basic emotion" theory (e.g., [5]), who assume a tight, biologically determined link between certain emotions and nonverbal behaviors, acknowledge that the information about emotions carried by these behaviors is limited to a small number of quality and intensity distinctions. And this is also what the empirical data suggest (e.g., [15], [20], [32], [38]). Nonverbal behaviors do not provide information about the subtler distinctions between human emotions, and may entirely miss emotions of low intensity. Most nonverbal behaviors (facial expression, tone of voice, physiological changes) do not, by themselves, provide information about the object of the emotion (what the person is happy about or afraid of), nor about the specific beliefs and desires that caused an emotion, or the specific action intentions that the emotion may have caused. However, these kinds of information are of crucial importance for being able to appropriately react to, and even to fully understand, another's emotions.

Even the limited amount of information about affect that is potentially available in nonverbal expression is often unavailable. The main reason for this is that emotions are not always mirrored in the nonverbal systems or not always in the same way. These problems can be well illustrated for facial expression, generally considered to be the best-discriminating nonverbal channel of emotion expression, and partly for this reason the nonverbal channel on which most AD research has focused (Calvo and D'Mello, this issue, and [38]). (Similar conclusions can be made for other nonverbal indicators of emotion, such as bodily changes and paralinguistic features of speech; see, e.g., [15], [17], [20], [32].) To make the problem concrete, consider emotional reactions in computer games (e.g., [14], [36]). Even in simple games, humans can experience many different emotions in quick succession and with changing objects. For example, they can experience surprise about unexpected turns in the game, fear of being attacked by an enemy, and relief about getting away; hope of winning points, happiness about having won points, and disappointment about not doing so; pride about a clever move, and self-anger about a stupid one. The majority of these emotions (hope, relief, disappointment, pride) have no, or at least no unique, facial expression. But even if one considers only emotions (or groups of emotions) for which distinctive facial expressions are thought to exist, such as joy and surprise [5], facial expressions are far from perfect emotion indicators.

The relatively best evidence for the coherence of facial expression and emotion exists for smiling and associated emotions such as happiness and amusement (e.g., [31]). However, even in this case, the correlation between mental state and expression is by no means perfect. Not only can smiling occur in the absence of an emotion; smiles often fail to occur despite the presence of an appropriate emotion. Considerable research indicates that this is not only, and not even mainly, due to the suppression of the tendency to smile, as Ekman's [5] theory of basic emotions would suggest. Rather, the very occurrence of the tendency to smile seems to depend on factors in addition to emotion, in particular the social context (the presence of others; see [11], [32], [35]). With respect to other emotions for which distinctive facial expressions have been claimed [5], such as surprise, disgust, and anger, the association between emotion and expression seems to be even weaker. We have studied this question in some depth for the case of surprise [22], [25]. Our main empirical finding can be summarized by saying that the facial display of surprise is typically *not* shown by surprised people. For example, [22] found that no more than 34 percent of the participants showed at least one component of the traditional surprise expression (eyebrow raising, eye widening, mouth opening/jaw drop) in response to an unexpected, surprising solution to a quiz item. Furthermore, most surprise expressions were one-component displays. Reisenzein et al. [25], who induced surprise by unannounced audiovisual changes and the unexpected appearance of the participants' own face as the last picture in a series of photographs that had to be rated, observed surprise displays in only 12 percent of the cases and none of them was a full 3-component display. This study also provided evidence that the failure to show a surprise expression was not due to insufficient intensity of surprise or the deliberate suppression of the facial expression, for even if these factors were controlled, the frequency of facial expressions did not increase. The empirical evidence on the coherence of emotions and facial expressions for yet other "basic" emotions, such as anger and disgust, is less abundant, but points to the same conclusion (e.g., [26]): Emotional states are very often not revealed in facial expressions; rather, this is the case only under special circumstances.

The empirical findings of moderate to low correlations between emotions and facial expressions receive additional, deductive support from theoretical, specifically evolutionary-psychological considerations, which suggest that emotions should be signaled to others, if at all, only in a selective fashion. The reason is that by giving away their emotions, people incur costs: They become more predictable and thus exploitable by others, and they may give away useful information about the environment (e.g., that something unexpected happened in the case of surprise) for free (e.g., [3], [6]). The (truthful) signaling of emotions to others is therefore a form of biological altruism that, like other altruistic behaviors, requires special evolutionary conditions for its emergence. Possible evolutionary scenarios are kin selection, reciprocal altruism, group selection [29], and costly signaling. With the possible exception of costly signaling (see [3]), all of these scenarios require that emotions are not signaled *indiscriminatively*, but are revealed *selectively* to suitable targets—be it close kin, partners in a cooperative relationship, or members of a group with which the sender identifies.

7 Overcoming the Limitations of Affect Inference from Nonverbal Behaviors

To overcome the limitations of affect inference from nonverbal behaviors, information about emotion-eliciting events, emotion-related goal-directed actions, and self-reports needs to be taken into account. Information about the events that occur to a person can be used to infer the person's emotions together with an emotion theory that connects events to beliefs, desires, and emotions (see [2] for an affective computing example; see also, [33]). Likewise, knowledge about the occurrence of emotion-related actions can be used, together with a theory of mind that connects actions to emotions, to infer the presence of emotions from actions. For example, if one knows that Jim helped Joe, who suffered some mishap, one can infer that Jim may have felt pity for Joe ([37]).

However, the most direct and richest source of information about emotions, and emotion-related mental states, are self-reports. Compared to nonverbal emotion expressions (considered as emotion communications), verbal reports have two decisive advantages [24]. First, being intentional actions (speech acts), verbal communications allow us to precisely target the communication of an emotion to the right person at the right time. Second, verbal communications allow us to transmit much more, as well as more precise, information about emotions than nonverbal signals: At least in principle, one can communicate whatever feature of emotion is available to consciousness and thus to introspection. Both the quality and intensity of emotions are more precisely represented in consciousness than they are reflected in nonverbal expression [22]. In addition, introspection can yield direct, unambiguous information about the situational and mental context of the emotion, including the emotion's object (the emotion-eliciting event), the beliefs and desires that caused the emotion, and the action tendencies that it may, in turn, have caused. This information makes emotion attributions more secure, as well as more precise (see [25]). Nor are verbal communications restricted to reporting the occurrence of an emotion or describing it; *any* thought relating to an emotion episode can be communicated, including the comparison of the emotion to remembered cases, reflections on the emotion's normative appropriateness, or recommendations on how to deal with the emotion.

Given these advantages of verbal communications about emotions, it is not surprising to learn that they are the main means of communicating emotions in everyday life [30] and psychotherapy [8], as well as the most common means to assess emotions in psychological research. In fact, emotion self-reports are typically used as the gold standard to validate other methods of emotion measurement. Although self-reports are also used for this purpose in AD research, they are rarely used as a separate source of information in applied affect detection. The main reasons for this seem to be two perceived disadvantages of emotion self-reports, compared to nonverbal indicators of emotion. First, the flow of social interaction or task engagement is interrupted by giving a self-report. However, note that people do this all the time to some degree in everyday life, for example when they ask others whether they like an idea or a drink or when they spontaneously voice that they are surprised, happy, or disappointed about something the other said or did. Thus, on second thought, self-reporting one's emotions is by no means as unnatural or rare as it may at first seem to be, at least in many situations.

The second perceived disadvantage of verbal emotion communication is that it is subject to voluntary control, and therefore to deliberate suppression and falsification ("social editing," Calvo and D'Mello, this issue). However, I am not convinced that this concern is a serious obstacle to the use of self-reports in human-computer interaction. First, social editing is to some degree also possible for nonverbal emotion indicators. Second, these biases are but one source of measurement error; nonverbal behaviors are subject to other errors that are often more severe. Third, the possible biases of self-report could ultimately be considered, and corrected for, in a comprehensive system of affect detection. Fourth, if the interaction of a human with an artificial agent (e.g., a virtual instructor or therapist) is of a kind that makes the truthful reporting of emotions a desirable goal for the human, psychological theory predicts that self-reports will be truthful. In contrast, if the context suggests to the human interactants that they are better off withholding or falsifying their emotion self-reports, the attempt to nonetheless extract from them information about their emotions by sophisticated AD systems that combine behavioral and physiological indicators, contextual knowledge and possibly even brain scans (see Calvo and D'Mello, this issue) is ethically questionable. Under such circumstances, affect detection can quickly turn into undesired affect extraction, and mindreading becomes mindspying. These problems are currently not acute, but will become important as soon as workable AD systems are implemented on a broader scale. But, even now, both affective computing researchers and psychologists need to be acutely aware of this and related ethical problems raised by the emerging field of automated affect detection (see also [28]).

ACKNOWLEDGMENTS

Many thanks to Jon Gratch for his helpful comments on a previous version of this article.

REFERENCES

[1] M.G. Armentano and A. Amandi, "Plan Recognition for Interface Agents. State of the Art," *Artificial Intelligence Rev.*, vol. 28, pp. 131-162, 2007.

[2] C. Conati, "Probabilistic Assessment of User's Emotions in Educational Games," *J. Applied Artificial Intelligence*, vol. 16, pp. 555-575, 2002.

[3] J.-L. Dessalles, *Why We Talk: The Evolutionary Origins of Language.* Oxford Univ. Press, 2007.

[4] M.J. Doherty, *Theory of Mind: How Children Understand Others' Thoughts and Feelings.* Psychology Press, 2009.

[5] P. Ekman, "An Argument for Basic Emotions," *Cognition and Emotion,* vol. 6, pp. 169-200, 1992.

[6] A. Fridlund, *Human Facial Expression. An Evolutionary View.* Academic Press, 1994.

[7] N.H. Frijda, *The Emotions.* Cambridge Univ. Press, 1986.

[8] S.R. Fussell, *The Verbal Communication of Emotion: Interdisciplinary Perspectives.* Lawrence Erlbaum, 2002.

[9] A.I. Goldman, *Simulating Minds: The Philosophy, Psychology and Neuroscience of Mindreading.* Oxford Univ. Press, 2006.

[10] F. Heider, *The Psychology of Interpersonal Relations.* Wiley, 1958.

[11] M. Holodynski, "The Miniaturization of Expression in the Development of Emotional Self-Regulation," *Developmental Psychology,* vol. 40, pp. 16-28, 2004.

[12] A. Kapoor, B. Burleson, and R. Picard, "Automatic Prediction of Frustration," *Int'l J. Human Computer Studies,* vol. 65, pp. 724-736, 2007.

[13] R.J. Kate and R.J. Mooney, "Probabilistic Abduction Using Markov Logic Networks," *Proc. IJCAI-'09 Workshop Plan, Activity, and Intent Recognition,* July 2009.

[14] A. Kappas and A. Pecchinenda, "Don't Wait for the Monsters to Get You: A Video Game task to Manipulate Appraisals in Real Time," *Cognition and Emotion,* vol. 13, pp. 119-124, 1999.

[15] J.T. Larsen, G.G. Berntson, K.M. Poehlmann, T.A. Ito, and J.T. Cacioppo, "The Psychophysiology of Emotion," *Handbook of Emotions,* M. Lewis, J.M. Haviland-Jones, and L.F. Barrett, eds., third ed., pp. 180-195, Guilford, 2008.

[16] B.F. Malle, *How the Mind Explains Behavior: Folk Explanations, Meaning, and Social Interaction.* MIT Press, 2004.

[17] I.B. Mauss, F.H. Wilhelm, and J.J. Gross, "Is There Less to Social Anxiety than Meets the Eye? Emotion Experience, Expression, and Bodily Eesponding," *Cognition and Emotion,* vol. 18, pp. 631-662, 2004.

[18] A. Ortony, G. Clore, and A. Collins, *The Cognitive Structure of Emotions.* Cambridge Univ. Press, 1988.

[19] R.W. Picard, *Affective Computing.* MIT Press, 1997.

[20] B. Parkinson, "What Holds Emotions Together? Meaning and Response Coordination," *Cognitive Systems Research,* vol. 10, pp. 31-47, 2009.

[21] D.V. Pynadath and S. Marsella, "PsychSim: Modeling Theory of Mind with Decision-Theoretic Agents" *Proc. 19th Int'l Joint Conf. Artificial Intelligence,* pp. 1181-1186, 2005.

[22] R. Reisenzein, "Exploring the Strength of Association between the Components of Emotion Syndromes: The Case of Surprise," *Cognition and Emotion,* vol. 14, pp. 1-38, 2000.

[23] R. Reisenzein, "Emotions as Metarepresentational States of Mind: Naturalizing the Belief-Desire Theory of Emotion," *Cognitive Systems Research,* vol. 10, pp. 6-20, 2009.

[24] R. Reisenzein and M. Junge, "Language and Emotion from the Perspective of the Computational Belief-Desire Theory of Emotion," *The Conceptualisation of Emotion*, Lodz Studies in Language, vol. 20, B. Lewandowska-Tomaszczyk and P. Wilson, eds., Peter Lang, 2010 (in press).

[25] R. Reisenzein, S. Bördgen, T. Holtbernd, and D. Matz, "Evidence for Strong Dissociation between Emotion and Facial Displays: The Case of Surprise," *J. Personality and Social Psychology,* vol. 91, pp. 295-315, 2006.

[26] R. Reisenzein, S. Bördgen, M. Studtmann, and H. Weber, "Spontaneous Facial Expressions of Emotion: Data on Surprise, Disgust, and Anger," *Proc. XXIX Int'l Congress Psychology,* July 2008.

[27] R. Reisenzein and U. Rudolph, "The Discovery of Common-Sense Psychology," *Social Psychology,* vol. 39, pp. 141-150, 2008.

[28] C. Reynolds and R. Picard, "Affective Sensors, Privacy, and Ethical Contracts," *Proc. Conf. Human Factors in Computing Systems,* pp. 1103-1106, 2004.

[29] P.J. Richerson and R. Boyd, *Not by Genes Alone. How Culture Transformed Human Evolution.* Univ. of Chicago Press, 2005.

[30] B. Rimé, "Emotions Elicit the Social Sharing of Emotion: Theory and Empirical Review," *Emotion Rev.,* vol. 1, pp. 60-85, 2009.

[31] W. Ruch, "Will the Real Relationship between Facial Expression and Affective Experience Please Stand Up: The Case of Exhilaration," *Cognition and Emotion,* vol. 9, pp. 33-58, 1995.

[32] J.A. Russell, J.-A. Bachorowski, and J.M. Fernández-Dols, "Facial and Vocal Expressions of Emotion," *Ann. Rev. Psychology,* vol. 54, pp. 329-349, 2003.

[33] M. Siemer and R. Reisenzein, "The Process of Emotion Inference," *Emotion,* vol. 7, pp. 1-20, 2007.

[34] P. Thagard, *Coherence in Thought and Action.* MIT Press, 2000.

[35] H. Wagner and V. Lee, "Facial Behavior Alone and in the Presence of Others," *The Social Context of Nonverbal Behavior,* P. Philippot, R.S. Feldman, and E.J. Coats, eds., pp. 262-286, Cambridge Univ. Press, 1999.

[36] N. Wang and S. Marsella, "Introducing EVG: An Emotion Evoking Game," *Proc. Sixth Int'l Conf. Intelligent Virtual Agents,* Aug. 2006.

[37] B. Weiner, *Social Motivation, Justice, and the Moral Emotions: An Attributional Approach.* Erlbaum, 2006.

[38] Z. Zeng, M. Pantic, G.I. Roisman, and T.S. Huang, "A Survey of Affect Recognition Methods: Audio, Visual, and Spontaneous Expressions," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 31, no. 1, pp. 39-58, Jan. 2009.

Rainer Reisenzein received the PhD degree in psychology from the University of Salzburg, Austria. After spending a postdoctoral research year at the University of California Los Angeles, he was an assistant professor at the Free University Berlin and then at the University of Bielefeld. Since 2002, he has been a professor of general psychology at the University of Greifswald, Germany. His research focuses on theoretical and empirical questions related to emotion and motivation and is marked by an interdisciplinary orientation toward, in particular, philosophy (philosophy of science and philosophy of mind) and cognitive science. He has published theoretical, review, and empirical articles and book chapters in English and German on diverse emotion-related topics, including emotional experience, emotion concepts, the facial expression of emotion, basic emotions, cognitive emotion theory, and the history of emotion. A special research focus is the phenomenon of surprise in humans, considered from the perspective of emotion theory. He is also a coauthor (with Wulf-Uwe Meyer und Achim Schützwohl) of a German three-volume introduction to the psychology of emotions. In 2009, he edited (together with the philosopher Sabine Döring) a special issue of *Emotion Review* on the topic of emotional experience. He is on the editorial board of the journals *Emotion*, *Cognition and Emotion*, and *Experimental Psychology*.