Introduction

Computation of emotions in man and machines

Peter Robinson[1],* and Rana el Kaliouby[2]

[1]*Computer Laboratory, University of Cambridge, William Gates Building, 15 J.J. Thomson Avenue, Cambridge CB3 0FD, UK*
[2]*MIT Media Laboratory, 20 Ames Street, Cambridge, MA 02139, USA*

The importance of emotional expression as part of human communication has been understood since Aristotle, and the subject has been explored scientifically since Charles Darwin and others in the nineteenth century. Advances in computer technology now allow machines to recognize and express emotions, paving the way for improved human–computer and human–human communications.

Recent advances in psychology have greatly improved our understanding of the role of affect in communication, perception, decision-making, attention and memory. At the same time, advances in technology mean that it is becoming possible for machines to sense, analyse and express emotions. We can now consider how these advances relate to each other and how they can be brought together to influence future research in perception, attention, learning, memory, communication, decision-making and other applications.

The computation of emotions includes both recognition and synthesis, using channels such as facial expressions, non-verbal aspects of speech, posture, gestures, physiology, brain imaging and general behaviour. The combination of new results in psychology with new techniques of computation is leading to new technologies with applications in commerce, education, entertainment, security, therapy and everyday life. However, there are important issues of privacy and personal expression that must also be considered.

Keywords: psychology; emotions; human–computer interactions; affective computing

1. INTRODUCTION

The year 2009 marks the bicentenary of Charles Darwin's birth. We remember him today for the publication of *On the Origin of Species* in 1859 and, perhaps rather less, for *The Descent of Man* in 1871 and *The Expression of the Emotions in Man and Animals* a year later. The last of these (Darwin 1872) was a remarkable book, selling over 5000 copies on the day of publication, and another 4000 in the following four months. More importantly, it remains completely relevant today. In the book, Darwin explores the expression of emotions not only by humans but also by cats, dogs, horses and many other animals. He observes that the expressions within a species are similar, but differ between species, and considers the question of whether this expression is learned or innate.

It is only much more recently, perhaps in the last 40 years, that Darwin's theories of universality have been scientifically tested and our understanding of the importance of emotions in human communication developed. Ekman (1973) provided strong evidence for universality in a book published to mark the centenary of *The Expression of the Emotions in Man and Animals*. Twenty-five years later, Rosalind Picard observed that the latest scientific findings indicated that emotions play an essential role in rational decision-making, perception, learning and a variety of other cognitive emotions (Picard 1997). She went on to argue that effective communication with computers also requires emotional intelligence—computers must have the ability to recognize and express emotions.

The study of affective computing has subsequently blossomed, with research groups around the world investigating different aspects of the field. This paper presents a summary of the achievements and challenges involved in affective computing.

This special issue of the *Philosophical Transactions* follows a two-day Discussion Meeting on the *Computation of Emotions in Man and Machines* held at the Royal Society in April 2009. The speakers represented a wide spectrum of academic disciplines from psychology through computer science to design, and their papers present recent advances in theories of emotion and affect, their embodiment in computational systems, the implications for general communications and broader applications.

Four interlinked themes will be discussed in the following four sections:

(i) underlying theories from psychology that provide the basis for studying emotions;

---

* Author for correspondence (peter.robinson@cl.cam.ac.uk).

(ii) recognizing emotions expressed by humans through facial expressions, gestures and body posture;

(iii) expressing emotions through the same channels in animated cartoons and by robots; and

(iv) bringing these technologies to bear on practical applications.

A common observation across all four themes of the meeting was that affect-measuring technologies are becoming robust and reliable. It is now realistic to use the models and sensors in other experiments, allowing real-world testing of theories about the mechanisms that underpin affect, social cognition, developmental trajectories and modes of derailed development such as autism spectrum disorders. Previously unanswered questions across a number of disciplines can now be addressed systematically. This is a very exciting prospect.

These new tools range from automated facial expression analysis and mobile biosensors to avatars, robots and immersive environments, which can be precisely manipulated, allowing human–machine and human–human interactions to be examined in controlled, yet naturalistic experimental settings. For example, Baylor (2009) and Pelachaud (2009) manipulate appearance and abilities of embodied conversational agents to test theories of persuasion and trust.

Cohn (Boker *et al.* 2009) describes a novel experimental paradigm that allows for the systematic testing of dyadic interactions, an area that was previously hard to undertake because of the range of parameters involved. Frith (2009) shows how the brain's responses to robots are fundamentally different from those to humans and challenges us to develop robots that can produce an affective imprint in humans. Slater's immersive environments (Slater 2009) and Breazeal's (2009) social robots shed some light on the role of embodiment in social cognition. Robots can also be used to explore different aspects of human development and are already being used to test theories such as mentalizing and mirroring.

Affect-measurement technologies are particularly exciting because they enable studies of affect to move from the laboratory to the real world. Baron-Cohen *et al.* (2009) consider affect–comprehension abilities in children with autism spectrum disorders in everyday social situations. Picard (2009) presents simple, wearable sensors that can be used to study physiological indicators in large populations over extended periods, enabling detailed examination of the interaction between everyday experiences and physiology.

We conclude by presenting some observations that arose during discussions at the meeting, which set an agenda for future research in the field.

2. UNDERSTANDING EMOTION

The first theme explores the grounding in psychology for studying human emotions.

Ekman (2009) starts this special issue with a highlight of Darwin's contributions to the study of emotion and facial expression, ranging from the

theory of discrete emotions through the importance of facial expressions to the universality of emotions. Ekman emphasizes Darwin's methodological contributions that shaped the field of psychology as we know it today. Darwin introduced the concept of a judgement study that remains the most widely used method in studying facial expression; he also had excellent observational skills, which are especially important even as we move towards a technology that is capable of doing some or all of the processing and analysis. Ekman points out several areas that Darwin did not consider in his research. In some of these areas substantial progress has been made, such as in the study of the dynamics of facial expressions and the study of deception. Other areas remain unanswered, for example we do not yet have a clear mapping from facial expressions to emotions, nor do we have a clear understanding between indicators and communicative signals. Throughout, Ekman reminds us that all scientific endeavours are ultimately about people and relationships.

Frith (2009) looks more closely at the role of facial expressions in social interactions, demonstrating how a behaviour, such as facial movement, can evolve into a complex and rich communication system. Frith covers topics such as imitation and emotional contagion to mental state attribution and empathy, as well as human–human communication and human–robot communication. While Darwin's methods and his followers are deeply rooted in psychology and observation, Frith borrows methods from neuroscience and brain imaging. Frith describes two parallel processes through which sensory signals, including facial expressions, are converted into behaviour; one is associated with consciousness while the other is not. Frith gives numerous examples where the unconscious processing of facial expressions and eye-gaze direction plays an important role in social interactions, including trust and friendliness. In the conscious route, senders and receivers are cognizant of their signals, and can voluntarily control the exchange of signals.

Scherer (2009) draws on extensive empirical evidence from psychology and neuroscience to demonstrate that emotions are emergent processes, best described using a dynamic computational architecture. In his 'component process model of emotion', emotion is a cultural and psychobiological adaptation mechanism that is built on emotion-antecedent appraisal on different levels of processing to empower individuals to react flexibly and dynamically to internal and external events. Scherer proposes the emergent computational modelling of emotion using nonlinear dynamic systems.

Following these papers focused primarily on the face, de Gelder (2009) challenges the dominance of the face in emotion research and presents 12 reasons for including bodily expression in affective neuroscience. De Gelder surveys an emerging body of research that demonstrates how bodily expressions are recognized as reliably as facial expressions even under conditions of limited attention and awareness.

De Gelder also reviews differences between the two channels of emotional expression and discusses why they are both important. De Gelder highlights the

importance of a multi-modal approach—a theme that recurs throughout this special issue. One important reason for studying bodily expressions is that it provides an excellent opportunity to validate human emotion theories, the majority of which are currently built on studies of facial expressions. Multi-modal studies will determine whether existing theories generalize beyond the face to other affective signals.

3. AUTOMATIC RECOGNITION OF EMOTIONS

The second area of discussion covers machine sensing and recognition of emotions. Humans communicate a wide range of affective states through a rich and varied collection of non-verbal channels, including the face, voice, gestures and body posture. This area of research draws on many disciplines, including machine learning, pattern recognition, signal processing, machine vision, cognitive science and psychology.

Cohn (Boker *et al*. 2009) presents an experiment that shows how people adapt their head movements and facial expressions in response to their partners in a social exchange to maintain equilibrium throughout the interaction. The experiment uses real-time automatic face tracking to drive a synthetic facial avatar in real time for communication in one direction and straightforward video conferencing for the return channel. Attenuating the avatar's head movements led to greater head movement by the conversational partner, and attenuated facial expressions led to increased head nodding both by naive participants and even by confederates. These results show that the dynamics of head movements in dyadic conversation are adapted to maintain symmetry.

This novel experimental paradigm is very promising and enables researchers to study dyadic and small-group interactions in a controlled way that was not previously possible, enabling the manipulation of real-time appearance of facial structure and expression as well as head movements.

Hess *et al*. (2009) review a growing body of research that demonstrates that the face is not simply an empty canvas, but that facial appearance interacts with facial movements to modify the perception of emotion expressions. Facial appearance communicates important information such as sex, age, ethnicity, personality and other characteristics, which can define a person and the social group to which he or she belongs. Through a systematic manipulation of each of these parameters, Hess shows that this information affects the perception of emotion expressions because it shapes the expectations of the person on the decoding or receiving end and biases their inference and attribution.

Hess's work is important because it argues that while facial-expression researchers often remove context to control for its effect, this is almost impossible to do with facial appearance. Her work has implications on how videos are labelled for emotional content and by whom, since social beliefs about individuals inform how and what we read in faces.

Pantic (2009) reviews progress with machine analysis of facial expressions over the past 10 years. Most of the work so far has focused on the automatic detection of units in Ekman's Facial Action Coding System (Ekman & Friesen 1978) or the detection of prototypic expressions of the basic emotions.

Pantic highlights several challenges that researchers in this area have begun to address. The majority of methods published to date are not applicable to real-life situations where changes in facial expressions are subtle, fleeting and certainly non-prototypic. As more interest builds in the potential applications of automated facial analysis, the focus of the research in the field is shifting to automatic recognition of spontaneous facial movements. Another challenge is that of dynamic modelling of facial expressions. Pantic cites the growing body of literature in cognitive science that demonstrates the importance of temporal information in facial expression and presents trends in dynamic modelling of facial expressions. Finally, Pantic poses one of the prominent challenges in this area of research: the availability of publicly available datasets that are labelled for facial action units and emotional expressions.

Cowie (2009) challenges many current assumptions in emotion research, focusing on the perception of emotion. He argues that affective computing and machine-based methods for detection of emotion have shifted the emphasis from analysing a single stimulus (often a static image or a short audio clip) to a more realistic, immersive approach to the study of emotion. He reminds us of how emotions continuously, consciously and more often sub-consciously influence the way in which we see the world and interact with others. He reviews the state of the research field representing emotion in speech, facial expressions, gestures and body postures. Cowie also highlights the importance of considering the context of the particular task at hand when inferring emotions.

4. EXPRESSION OF EMOTIONS

The third area of discussion explores the possibility of emotional expression by computer-controlled robots and by computer-generated cartoon avatars. The main question to be considered is how people react to these synthetic expressions. Can a computer-generated expression of emotion elicit the same empathic response in a person? What minimal set of displays is required to elicit the response? Is there any effect analogous to the 'uncanny valley' in computer animation (Mori 1970) that just makes the displays seem frightening rather than empathic?

Breazeal (2009) presents her work on robots as social learners. Her central thesis is that people are social, so robots should also be social; physical robots should interact with people in terms of their interests, beliefs and desires. Within this framework, cognition (understanding what is being displayed) is less important than emotion (its immediate effects on the other participant in a conversation). Social learning in humans is based on imitation, so Breazeal has taken the interaction of a carer with a child as a model for interactions between people and robots. Her robots, notably Kismet (Breazeal 2002), look and act like babies to encourage exaggerated expressions of emotion in people, which can easily be

interpreted and imitated. As a consequence, many of the machine-learning problems associated with more analytical approaches are avoided. In effect, the human and the robot collaborate in a process of teaching and learning.

This approach presents four main challenges. The overall goal is to establish a framework in which the social and communicative skills of the robot can be aligned with those of the human participant. Collaborating as a learner and a teacher establishes a mutually agreeable level of complexity in the communication. The second problem is to filter the salient data out of the stream of information made available by physical sensors. Third, the robot has to plan appropriate actions in response to its stimuli. Finally, the robot has to evaluate the effects of its responses to diagnose errors and improve performance.

Pelachaud (2009) presents her work on modelling multi-modal expression of emotion in a virtual agent. Embodied conversational agents express emotions through facial expressions, gaze, body movement and gesture, and voice. The focus of her paper is the expression of emotion in the face and gestures of a computer-generated cartoon avatar. This expression is dynamic, evolving in three ways: a spatial extent controlling the magnitude of the movement, a temporal extent controlling the durations of its onset, apex and offset, and power controlling the acceleration of the movement. The main challenge is that emotions in human–computer interactions are rarely discrete and rarely limited to the six basic emotions. A dimensional representation with interpolation between reference expressions is more appropriate.

There are three practical complications. These blends must still be calculated even when the component expressions are contradictory. One solution is to divide the face into eight areas and blend expressions separately in each area. While doing this, the dynamics of the expressions must be controlled to give realism and interest. This cannot be achieved by deterministic finite state machine controls; a degree of random variation is necessary. Finally, the expressions must be appropriate for the social context. This involves modelling social conventions and personalities. The relative social standing and politeness of the participants will affect their expressions. Embodied conversational agents implementing these ideas to exhibit a vocabulary of nine emotions achieve recognition rates between 37 and 93 per cent.

Slater (2009) presents his work on place illusion and plausibility in immersive virtual environments. Immersive virtual reality systems present users with a $360°$ projected image, viewed through shutter glasses so that the left and right eyes receive different views allowing simulation of depth in the pictures. This can be a compelling experience, with a strong sense of presence allowing initial disbelief to be suspended. It can also be trite and risible. What influences the different experiences? Indeed, how can a user's response be measured?

Slater's approach is to seek a 'response as if real' in terms of physiological, behavioural and cognitive responses. He identifies two separate factors that must both be present if the user's response is to be as real. First, the virtual environment must provide a *place illusion* (π), the sense of being in a real place. Second, the environment must provide a *plausibility illusion* (ψ), the sense that the scenario being presented is actually happening. These two factors are independent, and absence of either will result in the user's response not being as if real. Place illusion depends on the degree to which the sensory motor contingencies afforded by the virtual reality system permit perception by using the participant's whole body, much as they would in reality. For example, head tracking enables them to bend down and look underneath things, and events in peripheral vision lead to automatic head turns. Plausibility illusion depends on the extent to which the system can produce events that directly relate to the participant, the overall credibility of the scenario being depicted in comparison with expectations. These are subjective concepts and can only be measured through indirect assessments such as physiological measures, behavioural responses and questionnaires. Interestingly, both can be achieved when the virtual reality only modestly approximates to physical reality.

Baylor (2009) presents her work on visual presence and appearance in virtual agents. She investigates what makes a virtual agent (such as an anthropomorphic interface agent or even a simple pop-up window containing text) persuasive. The result is to influence the empathy and receptiveness of a human viewer more strongly, leading to more significant responses in terms of interest, attitudes and beliefs. Key factors include the agent's visibility, gender, race, age and even attractiveness and 'coolness'.

Initial studies show that a visible agent is more persuasive than a spoken message, which in turn is more persuasive than a simple textual message or voice alone. This effect is even significant when a computer program is delivering an error message. A message spoken by a visible agent was much more positively received than a text message; participants tended to blame themselves when a text box popped up, but blamed the computer when the error message was delivered by the visible agent. As with communications between people, the appearance of virtual agents is also important. Users receive a message more readily from an avatar of their own gender and ethnicity. However, these general rules break down in specific contexts. Middle-school students (male and female) were significantly more persuaded to consider Engineering as a subject for further study when persuaded by a female (as opposed to male) agent. In a broader experimental study, avatars with young, 'cool' appearances were more effective in persuading university freshmen to study Engineering (especially when attractive and female), again suggesting a peer group effect. However, older, 'uncool' avatars were almost as effective, suggesting that old stereotypes of engineers persist.

5. APPLICATIONS
The final theme considers a range of emerging application areas for affective computing. First, affect-sensing and affect-expressing technologies can

enhance people's ability to make sense of and communicate with others, especially those individuals who have difficulties in understanding, communicating and regulating their affective systems, such as those with autism spectrum disorders, affective mood disorders, anxiety, non-verbal learning disorders, visually impaired individuals and others. Second, affect-sensing and affect-expressing technologies can advance social–emotional intelligence in machines including conversational agents and sociable robots to improve people's experience with technology. Finally, as affect-measurement technologies continue to mature, researchers will be able to build technologies and models that facilitate real-world testing of contending theories of the mechanisms that underpin social cognition, developmental trajectories and modes of derailed development such as autism. This is particularly exciting because it means that studies of affect can be transferred from the laboratory to real-world situations.

Baron-Cohen *et al*. (2009) bring a clinical perspective to the role of emotional intelligence in everyday social interactions and communication by highlighting the difficulties that individuals on the autism spectrum face. Individuals with autism spectrum conditions (ASCs) have major difficulties in recognizing and responding to emotional and mental states in others' facial expressions. Baron-Cohen starts with an operational definition of empathy, specifies two components of empathy, emphasizes how such difficulties underlie the social and communication difficulties that form a core of the diagnosis of ASCs, and argues that at least one of these components are amenable to teaching.

Baron-Cohen has developed a series of short animated films designed to enhance emotion comprehension in children with ASCs, and describes a study that evaluates the efficacy of the animated series in improving emotion vocabulary and recognition in clinical trials. The results are very promising, with three levels of generalization including both similar stimuli and substantially different video-based stimuli. Baron-Cohen concludes with a discussion on the importance of designing for these populations, and putting into consideration specific strengths (e.g. hyper-systemizing and attention to detail) as well as sensitivities (e.g. atypical sensory processing).

Picard (2009) continues on the theme of technology that helps people sense, regulate and communicate affect, focusing particularly on autonomic nervous system (ANS) arousal. Picard starts with a case study of people who have difficulty in expressing their emotions, appearing to be calm and relaxed, yet experiencing extreme overload as measured by ANS activation. Affect-sensing technology can change this, by enabling people with such difficulties to visualize and share emotional information with others. Preliminary trials show that the technology helps conversational partners understand the difficulties experienced by people with ASCs, even when those difficulties cannot be articulated clearly.

Picard also highlights an especially important application of these technologies—enabling researchers to validate contending theories of affect and social

cognition. She presents a new methodological paradigm where very large sets of individualized data are processed using dynamic pattern analysis that looks across many variables. This enables researchers to conduct real-world measurements in the field that were previously restricted to the laboratory.

Höök (2009) presents affective applications as affective loop experiences. She emphasizes that emotions are processes, part and parcel of the interaction, and covering everyday bodily, cognitive and social experiences. Within this context, she presents several prototypes that are built around this theme with the goal of pulling the user into the interaction and engaging the user at the affective, cognitive and social level. Users of the system are active participants of the interaction and can choose to drive and shape in ways that are meaningful to them.

Finally, Gaver (2009) brings a design perspective to affective computing. His observations raise difficult questions, criticize bland assumptions by technologists and introduce a healthy note of scepticism. Gaver challenges the use of computational approaches to emotion particularly in design, drawing on a range of technical, cultural and aesthetic reasons. Gaver describes a prototype that he built, which uses sensor-based inferences to comment on domestic 'well-being' and shares how this prototype was a failure. As an alternative, Gaver advocates a broader view of interaction design in which open-ended designs serve as resources for individual appropriation, and suggests that emotional experiences become one of several outcomes of engaging with them.

Numerous challenges emerged from this final theme. How far are our technologies from really working in the field under limited control? How close are our models in terms of modelling the complexity and uncertainty inherent in emotional and social interactions? In the case where this technology is presented as assistive or augmentative, generalization is an important concern. How well do the acquired skills generalize to new settings?

6. DISCUSSION

Each of these four sessions was followed by a discussion that explored the broader issues raised in the formal presentations. These provide a useful framework for future research.

From considering the way in which people understand emotions, it was apparent that a richer classification than the six basic emotions would be required. There are arguments for both multi-dimensional continua and larger taxonomies of discrete emotions. In either case, the simultaneous expression of several emotions must be accommodated in both analysis and synthesis. Similarly, it is apparent that still pictures cannot convey a full range of emotions; the dynamics of expression are also important.

More subtly, it is necessary to distinguish emotions from moods. Emotions are relatively transient external expressions, but moods are sustained aspects of an individual's personality. For example, anger is an emotion, but hatred is a mood. Most of the work

presented was considering emotions, but there is perhaps a need to understand the underlying moods.

Any analysis of mental states cannot take sensor readings in isolation, but must consider their context. Interpretation depends on culture, the activity being undertaken and even the time of day. There is also the question of personal variation, not only in amplitude of expression but also in personal displays such as blushing or dilation of pupils.

Many of the analytical systems presented relied on machine-learning techniques, and there was considerable discussion of what forms of training data were appropriate. There was a strong preference for naturalistic data, evoked by a real task. However, labelling the resulting data is difficult and time consuming. It was also argued that the difference between a real and acted smile is more a matter of degree than a discrete choice, but there is also evidence that authentic expressions of emotion originate from different neural pathways. Nevertheless, the case for using acted training data remains, especially if performed by professionally trained method actors.

It is hard to identify the minimal set of competences required in a system that is to express emotions. An animated avatar is more persuasive than text, but the quality of animation does not appear to be significant. Simply knowing that a robot is not human influences our own perceptions and expectations. Speech is valuable and prosody is important but, again, high fidelity does not appear to be significant. People are aware of their own moods and use this to control their emotional displays; in the same way, displays by avatars and robots must be appropriate to the context and audience. A degree of ambiguity in the synthesized displays may add to their interest and realism.

If there is one lesson to be learned from the development of popular computer applications over the past 25 years, it is that the most successful emerged by accident. So it seems futile to try to predict future applications of affective computing. It is clear that emotional intelligence is of some value, so applications should not be completely decoupled from their users' feelings. On the other hand, measurement of a person's mental state is unlikely to be exact. After all, even people find it hard. There will be some intermediate point where a degree of emotional intelligence will give hints to a system rather than precise guidance, but there is no obvious application waiting to occupy the sweet spot.

Perhaps surprisingly, participants in the discussion were generally relaxed about the social implications of computer systems that monitor and manipulate people's emotions. However, we were reminded of Aldiss's 1961 novel, *The primal urge*. The premise is the invention of an emotion register, a small disc on an individual's forehead wired through to its wearer's hypothalamus, so that that the disc glows red when its wearer is attracted to another person. After some initial ructions, these prove acceptable and even popular. However, the closing paragraph of the novel reveals the possibility of a second disc that glows when its wearer lies . . . . We are still some way from creating such a simple technology that gives such a clear signal of emotions, but there are still ethical questions to be addressed.

7. CONCLUSION

The remaining papers in this special issue of the *Philosophical Transactions* record the presentations made over the two days of the meeting. The computation of emotions draws on insights from a wide variety of disciplines, and these papers span the whole range from the underlying psychology through computer science and machine learning to practical applications. The authors represent the leading authorities in their fields and their work is setting the agenda for future research across the whole breadth of the subject. We hope that this compilation sets forth a research agenda for inter- and multi-disciplinary research in emotion research.

From the discussion meeting, several challenges emerged. The first challenge is to develop applications that address the wide range of affective states that people experience, communicate and attribute to each other. This entails continuing to push affect-sensing technologies and models that can identify subtle nuances and intensities of a wide range of affective states, which consider the rich spectrum of modalities that humans employ to express themselves, and which consider the dynamics of these channels. Models should represent emotions as emergent processes, which change with internal and external contextual cues, including a person's past experiences, appraisal of current events and relationship with the interaction partner.

The second challenge is brought about by adopting an embodied approach to sensing people's state, where systems are not just passive sensors but are instrumental in driving human–computer and human–human interactions. Besides the challenges brought about by sensing people 'in the wild', there are challenges in implementing real-time, dynamic mechanisms and novel applications that learn with their users rather than simply being trained in advance. With the increasing use of robots and agents that will interact with people, there is growing interest in enabling machines to learn better by learning from people in a more natural, collaborative manner, and not just from people who program the computer or robot.

The third grand challenge we propose is moving towards new methodologies that enable not only large-scale sensing of people's individual affect-related data but also networks of data, such as face-to-face interactions of groups of people, or physiological sensing of large audiences. Within this framework, every instance of an interaction can be considered as an opportunity to learn about that person's specific knowledge and experiences.

REFERENCES

Aldiss, B. W. 1961 *The primal urge*. New York, NY: Ballantine Books.

Baron-Cohen, S., Golan, O. & Ashwin, E. 2009 Can emotion recognition be taught to children with autism spectrum conditions? *Phil. Trans. R. Soc. B* **364**, 3567–3574. (doi:10.1098/rstb.2009.0191)

Baylor, A. L. 2009 Promoting motivation with anthropomorphic interface agents: the importance of agent presence and appearance. *Phil. Trans. R. Soc. B* **364**, 3559–3565. (doi:10.1098/rstb.2009.0148)

Boker, S. M., Cohn, J. F., Theobald, B.-J., Matthews, I., Brick, T. R. & Spies, J. R. 2009 Effect of damping head movement and facial expression in dyadic conversation using real-time facial expression tracking and synthesized avatars. *Phil. Trans. R. Soc. B* **364**, 3485–3495. (doi:10. 1098/rstb.2009.0152)

Breazeal, C. 2002 *Designing sociable robots*. Cambridge, UK: MIT Press.

Breazeal, C. 2009 The role of expressive behavior for robots that learn from people. *Phil. Trans. R. Soc. B* **364**, 3527–3538. (doi:10.1098/rstb.2009.0157)

Cowie, R. 2009 Perceiving emotion: towards a realistic understanding of the task. *Phil. Trans. R. Soc. B* **364**, 3515–3525. (doi:10.1098/rstb.2009.0139)

Darwin, C. 1872 *The expression of the emotions in man and animals*. London, UK: John Murray.

de Gelder, B. 2009 Why bodies? Twelve reasons for including bodily expressions in affective neuroscience. *Phil. Trans. R. Soc. B* **364**, 3475–3484. (doi:10.1098/rstb.2009.0190)

Ekman, P. 1973 *Darwin and facial expression*. New York, NY: Academic Press.

Ekman, P. 2009 Darwin's contributions to our understanding of emotional expressions. *Phil. Trans. R. Soc. B* **364**, 3449–3451. (doi:10.1098/rstb.2009.0189)

Ekman, P. & Friesen, W. V. 1978 *Facial action coding system: a technique for the measurement of facial movement*. Palo Alto, CA: Consulting Psychologists Press.

Frith, C. 2009 The role of facial expressions in social interactions. *Phil. Trans. R. Soc. B* **364**, 3453–3458. (doi:10. 1098/rstb.2009.0142)

Gaver, W. 2009 Designing for emotion (among other things). *Phil. Trans. R. Soc. B* **364**, 3597–3604. (doi:10. 1098/rstb.2009.0153)

Hess, U., Adams, R. B. & Kleck, R. E. 2009 The face is not an empty canvas; how facial expressions interact with facial appearance. *Phil. Trans. R. Soc. B* **364**, 3497–3504. (doi:10.1098/rstb.2009.0165)

Höök, K. 2009 Affective loop experiences: designing for interactional embodiment. *Phil. Trans. R. Soc. B* **364**, 3585–3595. (doi:10.1098/rstb.2009.0202)

Mori, M. 1970 The uncanny valley. *Energy* **7**, 33–35.

Pantic, M. 2009 Machine analysis of facial behaviour: naturalistic and dynamic behaviour. *Phil. Trans. R. Soc. B* **364**, 3505–3513. (doi:10.1098/rstb.2009.0135)

Pelachaud, C. 2009 Modelling multimodal expression of emotion in a virtual agent. *Phil. Trans. R. Soc. B* **364**, 3539–3548. (doi:10.1098/rstb.2009.0186)

Picard, R. W. 1997 *Affective computing*. Cambridge, UK: MIT Press.

Picard, R. W. 2009 Future affective technology for autism and emotion communication. *Phil. Trans. R. Soc. B* **364**, 3575–3584. (doi:10.1098/rstb.2009.0143)

Scherer, K. R. 2009 Emotions are emergent processes: they require a dynamic computational architecture. *Phil. Trans. R. Soc. B* **364**, 3459–3474. (doi:10.1098/rstb.2009.0141)

Slater, M. 2009 Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Phil. Trans. R. Soc. B* **364**, 3549–3557. (doi:10. 1098/rstb.2009.0138)