

The field of Affective Computing: An interdisciplinary perspective

Jonathan Gratch, University of Southern California

Introduction

In psychology, the term *affect* refers to a broad range of psychological constructs associated with balanced personal reactions and feelings. This includes momentary emotions but also longer-term states such as moods or personality traits. *Affective Computing* is the interdisciplinary field of study concerned with recognizing, understanding, simulating and stimulating affective states in the design of computational systems (see Calvo, D'Mello, Gratch, & Kappas, 2015 for a broad overview). Research in the area is motivated by the fact that emotion pervades human life – emotions motivate human behavior, they promote social bonds between people and between people and artifacts, and emotional cues play an important role in forecasting human mental state and future actions. Technology is less efficient if it perturbs human emotions; more efficient if it engages with them productively; more attractive if it appeals to human emotions; and often it is primarily concerned with enabling humans to experience particular emotions (notably happiness).

Since the coining of the term by Picard in 1997, affective computing has emerged as a cohesive and increasingly impactful discipline spanning computer science, psychology, neuroscience, philosophy, art and industry. The journal which I founded in 2010 (IEEE's *Transactions on Affective Computing*) has one of the highest impact factors for journals in computer science (7.17 at the time of this writing). Affective computing papers and keynotes are routinely presented at the International Society for Research on Emotions and the Society for Affective Science, as well as the dedicated conference on the topic (the International Conference on Affective Computing and Intelligent Interaction). Besides this scientific impact, affective computing is beginning to transform society. Technology giants such as Apple, Amazon, Google and Facebook, as well as hundreds of smaller companies are deploying affective computing methods to predict or influence consumer behavior. Indeed, the AI Now Institute, an interdisciplinary research center dedicated to understanding the social implications of artificial intelligence, listed affective computing as its number one societal concern in their 2019 report (Crawford et al., 2019).

In this article, I provide a broad overview of the field. Although many people are most familiar with techniques for automatically recognizing emotion, I will emphasize that affective computing is a broad and vibrant field that not only recognizes affective states, but attempts to model and predict affective responses, and uses these models to generate life-like robots and digital characters, as well as to shape how people make decisions. The next section defines what the field means by the term *affect*. This is then followed by sections that discusses, in turn, how affective computing methods focus on the mind, the body, and social interaction. I end with a brief discussion of the ethical issues raised by the field.

What is affect?

Although emotion is the most salient concept associated with the term affect, the field of affective computing addresses a broad range of psychological constructs (Scherer, 2010). These concepts include *emotions* such as angry or sad (which tend to be strong and short and directed at a specific event – e.g., “I am angry *at* you”), *moods* such as cheerful or gloomy (which are less intense, longer lasting and not specifically associated with a triggering event – e.g., “I feel irritable”), *interpersonal stances* towards other, such as warm or cold, preferences and attitudes towards objects or ideas (e.g., “I love that painting”), and *affective dispositions* such as nervous or anxious (which can persist for years).

Each of these affective constructs are seen as possessing important substructure. For example, emotion is seen as involving several component mechanisms (Fontaine, Scherer, Roesch, & Ellsworth, 2007). These include (a) appraisal processes (which are involved in triggering an emotional response), (b) psychophysiological changes (such as increased heart rate or amygdala activation), (c) motor expressions (such as facial expressions, vocal changes and gestures), (d) action tendencies (such as preparation for fight versus flight), (e) subjective experiences (such as self-reported feelings of anger), and (f) emotion regulation / coping processes (such as suppression or reappraisal).

Much of the recent controversies in emotion research surround the linkages between these different components. For example, Ekman’s basic emotion theory argues these components are tightly linked and essentially act as a single circuit (Ekman, 1992). The implication of this view is that recognition of behavior in one component is highly diagnostic of other components, and in particular, that felt emotions can be readily predicted from surface cues such as facial expressions. However, the emerging consensus in emotion research is that these components are only loosely connected. Thus, predicting felt emotions from facial expressions alone is difficult, if not impossible, without bringing in additional information such as details about the setting in which the expression was evoked (Barrett, Adolphs, Marsella, Martinez, & Pollak, 2019; Crivelli & Fridlund, 2018).

Other controversies surround the appropriate labels to distinguish the different states of these components. For example, emotion in particular has been argued to be best described by basic emotion categories such as hope and fear (Ekman, 1992), or richer basic categories – e.g., Ortony and colleagues distinguish twenty-two emotion categories (Ortony, Clore, & Collins, 1988) – or continuous dimensions such as valance and arousal (Russell, 1980) or evaluation, potency and activity (Osgood, May, Miron, & Miron, 1975). Further, emotion research has tended to overemphasize distinctions amongst negative states but recent work has begun to emphasize distinctions between positive affective states (Shiota et al., 2014).

These different senses of affect can create confusion and potentially lead techniques to be used in inappropriate ways. Unlike some scientific disciplines, affective computing uses terms, such as emotion, that have intuitive meanings in popular culture that differ from their scientific usage. This confusion can be amplified by marketing hype by commercial companies that sell these techniques. For example, Amazon claims its Rekognition software can “detect emotions such as happy or sad”¹ and that this can be effective for “determining if a movie is subjectively funny.”² The commonsense interpretation of these statements is that the software predicts what people would say they feel about a situation. Yet deep in the software documentation one can find that the software detects “emotions

¹ <https://aws.amazon.com/getting-started/hands-on/detect-analyze-compare-faces-rekognition/>

² https://aws.amazon.com/deeplens/community-projects/Facial_Emotion/

that appear to be expressed on the face.”³ In other words, the software was trained to predict what a 3rd party would think a person is expressing, when looking at an image of their face (e.g., if are they smiling). Unfortunately, the association between expression and self-reported feeling is now known to be quite weak. For example, people smile for many reasons including out of politeness (Ambadar, Cohn, & Reed, 2009), in response to surprising events (Lei & Gratch, 2019) or even frustration (Hoque & Picard, 2011). Seeing that a person smiles when they watch a movie provides some evidence that they find it funny, but such interpretations must be taken with care and are best combined with other evidence and contextual factors. This is particularly true when expressions are used for societally important issues like hiring (Harwell, 2019) or law enforcement (Rose, 2019).⁴

Mind

Affective computing techniques focus on far more than recognizing affective signals. Affects are closely intertwined with how people make decisions (Damásio, 1994; Loewenstein & Lerner, 2003). For example, people make riskier decisions in a happier mood and choices are often made based on the anticipated emotions that they will bring about. By modeling these influences, affective computing techniques can, for example, predict the emotions people will feel (Gratch & Marsella, 2005) and how these emotions influence decision-making (Busemeyer, Dimperio, & Jessup, 2007).

Most models of affect and decision-making build upon appraisal theory, one of the most widely-accepted theories on the antecedents of emotion (e.g., see Scherer, 2005). There are several variations of appraisal theory but they share many of the basic processes illustrated in Figure 1. Affect is argued to arise from an appraisal process that compares events in the world (or imagined events in the mind) with an organism’s beliefs and goals. This comparison involves several facets, often called *appraisal variables*, including if the event facilitates or hinders the organism’s goals, if the event was expected, the level of control over the event, and who is to blame for the good or bad consequences. Different patterns of appraisal will tend to elicit specific affective states (e.g., fear arises when bad circumstances seem uncontrollable, whereas anger arises when control seems high) and prepare the organism to respond in adaptive ways. For example, when angry, blood tends to flow to the arms, preparing a person to attack (Lerner & Tiedens, 2006). These evoked affective states serve as control signals that guide subsequent behavior. They can motivate action in the world (for example, expending effort to remove an obstacle to a goal), they can lead the individual to reevaluate their beliefs and goals (e.g., giving up on a goal that is blocked) or can lead the individual to attempt to suppress these feelings.

Several computational methods have been proposed to model these processes. A central challenge is how to derive appraisal variables in a general and domain-independent fashion. For example, the EMA model (Gratch & Marsella, 2004; Marsella & Gratch, 2009) uses AI planning methods to represent situations and derives appraisal variables from inferences over planning representations (e.g., is the establishment of a goal threatened by some other action). As another example, Broekens and colleagues use reinforcement learning to derive appraisals from a set of state-transition and reward functions (Moerland, Broekens, & Jonker, 2016). Other models propose extensions to classical decision

³ https://docs.aws.amazon.com/rekognition/latest/dg/API_Emotion.html

⁴ Note that companies like Amazon sometimes include disclaimers though these can be difficult to find. In particular, Amazon includes caveats for the use of its software in situations involving public safety including “Don’t use Amazon Rekognition responses to make autonomous decisions for scenarios that require analysis by a human”, though this caveat is made for face recognition, not emotion recognition (<https://docs.aws.amazon.com/rekognition/latest/dg/considerations-public-safety-use-cases.html>)

theory. For example, Meller's decision-affect theory modifies the standard expected utility calculation by incorporating appraisals of surprise, disappointment and regret (Mellers, Schwartz, Ho, & Ritov, 1997). See (Marsella, Gratch, & Petta, 2010) for a review of several computational appraisal models.

Computational appraisal models can be used to make predictions about how people will feel (Elliott, 1992; Mao & Gratch, 2012), what decisions will be made (Mellers et al., 1997), and how people cope with strong emotions, for example through resignation or wishful thinking (Marsella & Gratch, 2003). Given they capture "typical" human behavior, they can help identify clinical conditions where people depart from normal decision-making (Yechiam, Busemeyer, Stout, & Bechara, 2005). Given that they can predict how situations typically unfold (Robinson, Smith-Lovin, & Wisecup, 2006), they have been used to create or analyze stories (Bergstrand & Jasper, 2018) and drive the behavior of lifelike robots (Kim & Kwon, 2010) and virtual characters (Neal Reilly, 1996; Traum, Swartout, Marsella, & Gratch, 2005).

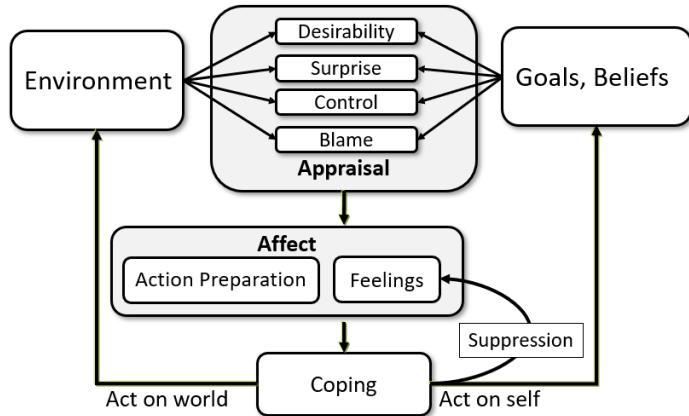


Figure 1: An illustration of appraisal theory Body

Affects not only impact decision-making, but are associated with altered behavior across a wide range of physiological and motor processes. Modeling and recognizing these changes can be important in themselves (e.g., predicting long-term bodily consequences of stress). But in that affect involves multiple components (i.e., appraisal, psychophysiological change, motor expressions, action tendencies, subjective feelings, and emotion regulation), changes in one component can allow inferences about changes in another. For example, recognizing sympathetic versus parasympathetic nervous system activation may help predict if someone is feeling a strong emotion (Bradley, Miccoli, Escrig, & Lang, 2008) and modeling how decisions are shaped by physiological needs, like hunger and thirst, can give insights into social and individual decision-making (Dancy, 2019) .

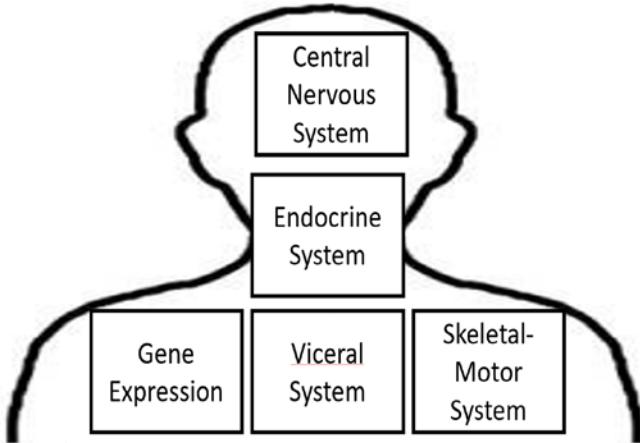


Figure 2: Major body systems impacted by affect

Figure 2 illustrates the different physiological and motor processes that are impacted by affective states. With the exception of gene expression and the endocrine system, each of these elements has been a focus of affective computing techniques. Much of the focus of affective computing has been on recognizing and characterizing processes within each of these elements, but work is also advancing on modeling these bodily processes, which in turn, can benefit both recognition and generation techniques.

Several scholars have emphasized the importance of developing computational models of bodily processes (e.g., Fellous, Armony, & LeDoux, 2002; Scherer, 2009). Much of the work in this area has focused on detailed models of specific bodily elements. For example, models have been produced of how emotion impacts the amygdala (Morén & Balkenius, 2000) or pupil dilation (Johansson & Balkenius, 2018), and some of these models have even inspired the design of practical products, such as power control systems (Rahman, Milasi, Lucas, Araabi, & Radwan, 2008). But affective computing research is beginning to explore comprehensive models of human physiology. For example, HumMod simulates the interaction of 5000 physiological variables including the visceral, endocrine, and skeletal-motor systems (Hester et al., 2011), and ACT-R/Φ combines physiological and cognitive models to better explain how stress shapes decision-making (Dancy, Ritter, Berry, & Klein, 2015).

The importance of models notwithstanding, most affective computing research has focused on how to recognize activation within each of these physiological elements. Work on the central nervous system has emphasized detection of affective states from electroencephalogram (EEG) signals (see Kim, Kim, Oh, & Kim, 2013; Mühl, Heylen, & Nijholt, 2015 for recent reviews), although recent work is beginning to look at other measures including fMRI. For example, Kragel and colleagues trained a classifier to predict the emotions evoked from video clips and argued that the learned representations correspond to brain activation in human subjects watching the same videos (Kragel, Reddan, LaBar, & Wager, 2019). The visceral system has been explored with a wide range of measures including skin-conductance (Sano & Picard, 2013), cardiography (Khooshabeh et al., 2013; Nardelli, Valenza, Greco, Lanata, & Scilingo, 2015), pupil dilation (Partala & Surakka, 2003). Work on the skeletal-motor system has emphasized facial expression (Cohn & De la Torre, 2015) and bodily gestures and postures (Bianchi-Berthouze & Kleinsmith, 2015).

Although less studied within affective computing, affective states are also manifest in changes in hormones and even result in genetic changes. For example, long-term stress can be measured by the

endocrine hormone cortisol (Hellhammer, Wüst, & Kudielka, 2009) and positive emotions can impact gene expression in the hypothalamus (Hori et al., 2009). Some work has begun to explore real-time direct measurement of such signals, including gene expression in mice (Ono, Honma, & Honma, 2015) and human cortisol that can be measured with a smartphone (Choi et al., 2014), but these measures are still invasive and impractical for many settings. One promising direction is to indirectly detect physiological changes using a less invasive measure. For example, Scherer and colleagues explored the relationship between physiological changes and changes in the voice (Neubauer et al., 2017) and Abdelnasser and colleagues showed that respiration rate could be measured non-invasively using just the WiFi in the home (Abdelnasser, Harras, & Youssef, 2015).

In addition of recognition, affective computing considers how to automatically generate affective manifestations. Indeed, there is an increasingly close connection between recognition and generation as models for recognition can sometimes be converted into generative models and both recognition and generation share common intermediate feature representations. Generation work includes how to simulate the influences of affect on speech, both in terms of prosody (Burkhardt & Campbell, 2015) and lexical content (Wang & Wan, 2018), posture and gestures (Calvo, D'Mello, et al., 2015) and facial display (Ochs, Niewiadomski, & Pelachaud, 2015).

Society

Affective states are manifest through bodily and behavior changes, and these manifestations can alter the behavior of others. For example, people are more willing to make concessions to an angry negotiator (van Kleef, de Dreu, & Manstead, 2004), even when those expressions are generated by a machine (de Melo, Carnevale, & Gratch, 2011). These signals can even propagate through social networks like Facebook (Kramer, Guillory, & Hancock, 2014). More fundamentally, many affective states are inherently social (e.g., guilt or envy), serve important social functions (e.g., anger is a signal that a relationship is broken and needs to be repaired), and are “co-constructed” through the back and forth between individuals: for example, if one person leans forward and the other leans back, this may be interpreted as fear whereas if one person leans forward and the other leans forward as well, this may be interpreted as anger (see Parkinson, Fischer, & Manstead, 2005).

There is some disagreement within emotion research concerning the mechanisms by which these manifestations shape social behavior. Research on emotion contagion argues people essentially catch emotions from others, often below the level of conscious awareness (Hatfield, Cacioppo, & Rapson, 1994). Some affective computing work follows this view, for example by modeling how fear propagates through crowd just as heat propagates (Tsai, Bowring, Marsella, & Tambe, 2013) and Zhang and colleagues apply a contagion model when proposing a method for calming crowds in an evacuation (Zhang, Lu, & Liu, 2018).

The more prominent view is these manifestations serve as information. For example, van Kleef's Emotions as Social Information (EASI) model argues that people make inferences about the likely mental state and future behavior of a person based on the emotional reactions they show to specific events. The most detailed of these models build off of the appraisal theory illustrated in Figure 1. In what the authors refer to as “reverse appraisal”, de Melo and colleagues argue that when observers see a specific expression following an event, they use this as information about how the producer is appraising the situation, and by extension, the goals the person holds (de Melo, Carnevale, Read, & Gratch, 2014). For example, in the prisoner's dilemma game, if my opponent shows guilt after exploiting me, this suggests that the exploitation was accidental and I should give them another chance before punishing them. In

contrast, if my opponent smiles after exploiting me, this suggests they will exploit me again if given another chance (see also Shore, Rychlowska, van der Schalk, Parkinson, & Manstead, 2018).

Reverse appraisal provides a useful framework for helping machines generate persuasive emotional signals. For example, de Melo showed how to influence people's willingness to cooperate with a virtual agent based on a pattern of emotional signals that was predicted by reverse appraisal (de Melo, Carnevale, & Gratch, 2010). We have also shown that reverse appraisal can help agents effectively lie in negotiations – helping to exploit an opponent while convincing them that they were treated fairly (Gratch, Nazari, & Johnson, 2016). Other work has shown that other factors, such as culture, can shape how people interpret the emotional signals of machines (de Melo & Terada, 2019).

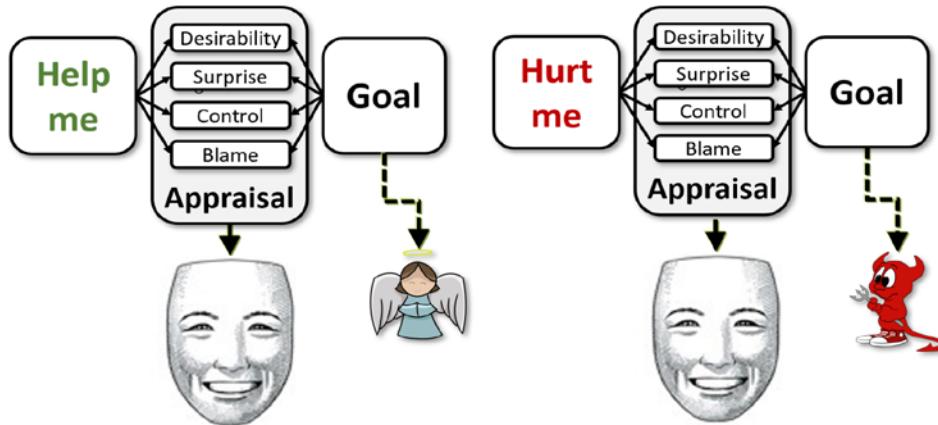


Figure 3: Illustration of reverse appraisal. Smiles following helping promote pro-social inferences whereas smiles following harming promote anti-social inferences.

Up to now we have focused on how artifacts can act in human-like ways and use human-like signals to alter social behavior. Art, music, architecture, even product design can also be seen as a means to evoke affective responses, and affective computing has begun to impact these creative fields. In Japan, Kansei engineering seeks to aid in the development of products or services that evoke specific customer feelings and needs (Nagamachi, 1995). In general, the factors that elicit emotions from an object may depend on quite different factors than how they are evoked by another person. For example, Menninghaus' research into "aesthetic emotions" argues that art tends to evoke states rarely considered by human-centric emotion research, including enchantment, nostalgia, unease and insight (Menninghaus et al., 2019). Juslin's theory of music emotions emphasizes such factors as rhythmic entrainment and brain stem reflexes (Juslin, 2013). To the extent that such aesthetic states could be modeled and predicted, it could have important practical implications, such as improved music recommendation (Soleymani, Aljanaki, Wiering, & Veltkamp, 2015) or in the automatic generation of emotionally-evocative objects (e.g., Utz & DiPaola, 2020).

Bias and ethics

In that affective computing seeks to recognize affective states, uses those recognitions to make inferences (e.g., is this person trustworthy? Do they like my product?) and exploits affective signals to shape human decisions, questions of ethics and bias are rightly coming to the forefront. As alluded to in the introduction, AI Now's 2019 report had as its number one recommendation that "Regulators should

ban the use of affect recognition in important decisions that impact people's lives and access to opportunities" (Crawford et al., 2019, pg 6). There are many problems with that recommendation. In particular, although the report calls for the ban of affect recognition in general, they base this on quite narrow evidence: i.e., evidence that machines cannot predict basic emotion categories (e.g., hope, joy, fear..) from decontextualized images of the face (i.e., Barrett et al., 2019). Nonetheless, there is good reason to be cautious in the use of this technology and any application should be evidence-based. As noted in Section 2, there is often confusion around the meaning of affect terms, and naïve users of the technology can easily use the techniques in inappropriate ways. Further, affective computing is subject to many of issues of algorithmic bias faced by other areas of AI. For example, Figure 4 illustrates some examples from our own lab of common problems faced by facial expression recognition technology. For the woman on the left, the software believes her name tag represents the only face in the image and it is attempting to make sense of that as a facial expression. The second image illustrates that background items can prove attractive distractors – in this case the agent is fixated on Abraham Lincoln's face in the background. More systematically, the chart on the right shows the accuracy of a commercial face detection algorithm we use in our lab – it is much less accurate detecting black faces. In general, affective recognition accuracy can be impacted by a variety of contextual factors. Recognized expressions change systematically as people look away from the camera – for example, as people look down, algorithms are more convinced they are smiling (Kappas, Hess, Barr, & Kleck, 1994) – or by the location of lights in the environment (Stratou, Ghosh, Debevec, & Morency, 2012) or when they are speaking (Kim & Mower Provost, 2014).

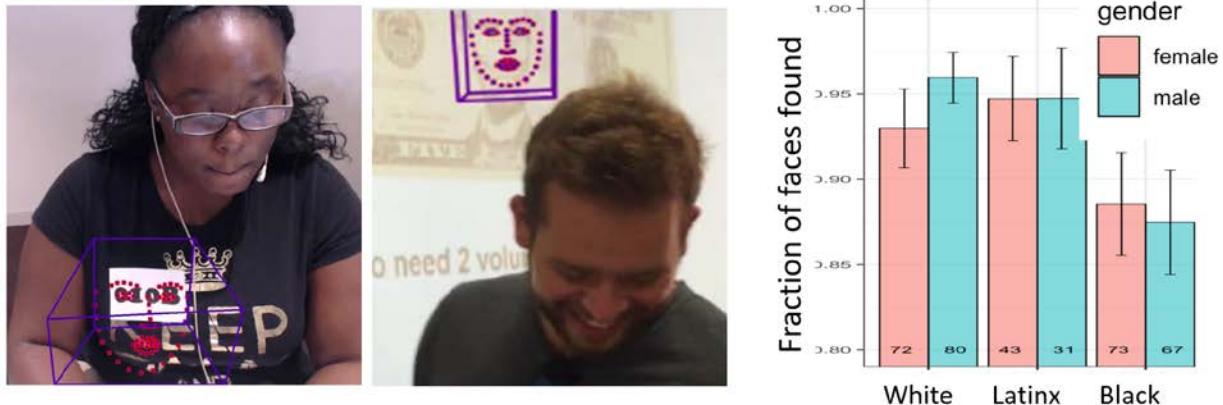


Figure 4: Illustrations of sensitivity of affective computing methods to race and background clutter.

Accounting for the impact of context enhance the accuracy of affective-computing models, but accuracy is not the only ethical concern. As noted the previous section, algorithms can use emotional signals to shape social outcomes. Even when the goals of these systems seem benign – e.g., helping people take their medicine (Bickmore, Gruber, & Picard, 2005) or disclose symptoms of mental illness (Lucas, Gratch, King, & Morency, 2014), the use of emotional manipulation may be problematic. The problems only magnify when we consider applications where agents manipulate for their own benefit at the cost of the user (Gratch, 2020). For discussions of ethics in affective computing, see (Beavers & Slattery, 2017; Cowie, 2015; Reynolds & Picard, 2004).

Acknowledgements

Research was sponsored by the Army Research Office under Cooperative Agreement Number W911NF-20-2-0053, the Air Force Office of Scientific Research under grant FA9550-18-1-0182, and the National Science foundation under grant 1822876. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

Abdelnasser, H., Harras, K. A., & Youssef, M. (2015). *UbiBreathe: A ubiquitous non-invasive WiFi-based breathing estimator*. Paper presented at the Proceedings of the 16th ACM International Symposium on Mobile Ad Hoc Networking and Computing.

Ambadar, Z., Cohn, J. F., & Reed, L. I. (2009). All smiles are not created equal: Morphology and timing of smiles perceived as amused, polite, and embarrassed/nervous. *Journal of nonverbal behavior*, 33(1), 17-34.

Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychological Science in the Public Interest*, 20(1), 1-68. doi:10.1177/1529100619832930

Beavers, A. F., & Slattery, J. P. (2017). On the moral implications and restrictions surrounding affective computing *Emotions and Affect in Human Factors and Human-Computer Interaction* (pp. 143-161): Elsevier.

Bergstrand, K., & Jasper, J. M. (2018). Villains, Victims, and Heroes in Character Theory and Affect Control Theory. *Social Psychology Quarterly*, 81(3), 228-247. doi:10.1177/0190272518781050

Bianchi-Berthouze, N., & Kleinsmith, A. (2015). Automatic recognition of affective body expressions *The Oxford Handbook of Affective Computing*.

Bickmore, T., Gruber, A., & Picard, R. (2005). Establishing the computer-patient working alliance in automated health behavior change interventions. *Patient Education and Counseling*, 59(1), 21-30.

Bradley, M. M., Miccoli, L., Escrig, M. A., & Lang, P. J. (2008). The pupil as a measure of emotional arousal and autonomic activation. *Psychophysiology*, 45(4), 602-607. doi:10.1111/j.1469-8986.2008.00654.x

Burkhardt, F., & Campbell, N. (2015). Emotional speech synthesis *The Oxford Handbook of Affective Computing* (pp. 286): Oxford University Press New York.

Busemeyer, J. R., Dimperio, E., & Jessup, R. K. (2007). Integrating emotional processing into decision-making models. In W. D. Grey (Ed.), *Integrated models of cognitive systems*. Oxford: Oxford University Press.

Calvo, R., D'Mello, S., Gratch, J., Kappas, A., Lhommet, M., & Marsella, S. (2015). Expressing emotion through posture and gesture: Oxford University Press.

Calvo, R. A., D'Mello, S., Gratch, J., & Kappas, A. (2015). *The Oxford handbook of affective computing*: Oxford University Press, USA.

Choi, S., Kim, S., Yang, J.-S., Lee, J.-H., Joo, C., & Jung, H.-I. (2014). Real-time measurement of human salivary cortisol for the assessment of psychological stress using a smartphone. *Sensing and Bio-Sensing Research*, 2, 8-11. doi:<https://doi.org/10.1016/j.sbsr.2014.08.001>

Cohn, J. F., & De la Torre, F. (2015). Automated face analysis for affective computing.

Cowie, R. (2015). Ethical issues in affective computing. *The Oxford Handbook of Affective Computing*, 334.

Crawford, K., Dobbe, R., Dryer, T., Fried, G., Green, B., Kaziunas, E., . . . Sánchez, A. N. (2019). AI now 2019 report. *New York, NY: AI Now Institute*.

Crivelli, C., & Fridlund, A. J. (2018). Facial Displays Are Tools for Social Influence. *Trends in Cognitive Sciences*, 22(5), 388-399. doi:<https://doi.org/10.1016/j.tics.2018.02.006>

Damásio, A. R. (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Avon Books.

Dancy, C. L. (2019). A hybrid cognitive architecture with primal affect and physiology. *IEEE Transactions on Affective Computing*, 1-1. doi:10.1109/TAFFC.2019.2906162

Dancy, C. L., Ritter, F. E., Berry, K. A., & Klein, L. C. (2015). Using a cognitive architecture with a physiological substrate to represent effects of a psychological stressor on cognition. *Computational and Mathematical Organization Theory*, 21(1), 90-114.

de Melo, C., Carnevale, P., & Gratch, J. (2010). The Influence of Emotions in Embodied Agents on Human Decision-Making. In J. Allbeck, N. Badler, T. Bickmore, C. Pelachaud, & A. Safanova (Eds.), *Intelligent Virtual Agents* (Vol. 6356, pp. 357-370): Springer Berlin Heidelberg.

de Melo, C., Carnevale, P. J., & Gratch, J. (2011). *The Effect of Expression of Anger and Happiness in Computer Agents on Negotiations with Humans*. Paper presented at the the Tenth International Conference on Autonomous Agents and Multiagent Systems, Taipai, Taiwan.

de Melo, C., Carnevale, P. J., Read, S. J., & Gratch, J. (2014). Reading people's minds from emotion expressions in interdependent decision making. *Journal of Personality and Social Psychology*, 106(1), 73-88. doi:10.1037/a0034251

de Melo, C. M., & Terada, K. (2019). Cooperation with autonomous machines through culture and emotion. *PLoS One*, 14(11), e0224758.

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

Elliott, C. (1992). *The affective reasoner: A process model of emotions in a multi-agent system*. (PhD Ph.D Dissertation), Northwestern University, Northwestern, IL.

Fellous, J., Armony, J., & LeDoux, J. (2002). Emotional Circuits and Computational Neuroscience. The Handbook of Brain Theory and Neural Networks: Cambridge, MA: The MIT Press.

Fontaine, J. R. J., Scherer, K. R., Roesch, E. B., & Ellsworth, P. C. (2007). The World of Emotions is not Two-Dimensional. *Psychological Science*, 18(12), 1050-1057. doi:10.1111/j.1467-9280.2007.02024.x

Gratch, J. (2020). *The promise and peril of automated negotiators*. Working Paper.

Gratch, J., & Marsella, S. (2004). A domain independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 5(4), 269-306.

Gratch, J., & Marsella, S. (2005). Evaluating a computational model of emotion. *Journal of Autonomous Agents and Multiagent Systems*, 11(1), 23-43.

Gratch, J., Nazari, Z., & Johnson, E. (2016). *The Misrepresentation Game: How to win at negotiation while seeming like a nice guy*. Paper presented at the International Conference on Autonomous Agents and Multiagent Systems, Singapore.

Harwell, D. (2019, Nov 6, 2019). A face-scanning algorithm increasingly decides whether you deserve the job. *Washington Post*. Retrieved from <https://www.washingtonpost.com/technology/2019/10/22/ai-hiring-face-scanning-algorithm-increasingly-decides-whether-you-deserve-job/>

Hatfield, E., Cacioppo, J. T., & Rapson, R. L. (Eds.). (1994). *Emotional Contagion*. Cambridge: Cambridge University Press.

Hellhammer, D. H., Wüst, S., & Kudielka, B. M. (2009). Salivary cortisol as a biomarker in stress research. *Psychoneuroendocrinology*, 34(2), 163-171. doi:<https://doi.org/10.1016/j.psyneuen.2008.10.026>

Hester, R., Brown, A., Husband, L., Iliescu, R., Pruett, W. A., Summers, R. L., & Coleman, T. (2011). HumMod: a modeling environment for the simulation of integrative human physiology. *Frontiers in physiology*, 2, 12.

Hoque, M., & Picard, R. W. (2011, 21-25 March 2011). *Acted vs. natural frustration and delight: Many people smile in natural frustration*. Paper presented at the Face and Gesture 2011.

Hori, M., Hayashi, T., Nakagawa, Y., Sakamoto, S., Urayama, O., & Murakami, K. (2009). Positive emotion-specific changes in the gene expression profile of tickled rats. *Mol Med Rep*, 2(2), 157-161. doi:10.3892/mmr_00000077

Johansson, B., & Balkenius, C. (2018). A computational model of pupil dilation. *Connection Science*, 30(1), 5-19. doi:10.1080/09540091.2016.1271401

Juslin, P. N. (2013). From everyday emotions to aesthetic emotions: Towards a unified theory of musical emotions. *Physics of Life Reviews*, 10(3), 235-266. doi:<https://doi.org/10.1016/j.plrev.2013.05.008>

Kappas, A., Hess, U., Barr, C. L., & Kleck, R. E. (1994). Angle of regard: The effect of vertical viewing angle on the perception of facial expressions. *Journal of nonverbal behavior*, 18(4), 263-280. doi:10.1007/bf02172289

Khooshabeh, P., de Melo, C., Volkman, B., Gratch, J., Blascovich, J., & Carnevale, P. J. (2013). *Negotiation strategies with incongruent facial expressions of emotion cause cardiovascular threat*. Paper presented at the 35th annual meeting of the Cognitive Science Society, Berlin, Germany.

Kim, H.-R., & Kwon, D.-S. (2010). Computational Model of Emotion Generation for Human–Robot Interaction Based on the Cognitive Appraisal Theory. *Journal of Intelligent & Robotic Systems*, 60(2), 263-283. doi:10.1007/s10846-010-9418-7

Kim, M.-K., Kim, M., Oh, E., & Kim, S.-P. (2013). A Review on the Computational Methods for Emotional State Estimation from the Human EEG. *Computational and Mathematical Methods in Medicine*, 2013, 573734. doi:10.1155/2013/573734

Kim, Y., & Mower Provost, E. (2014). *Say cheese vs. smile: Reducing speech-related variability for facial emotion recognition*. Paper presented at the Proceedings of the 22nd ACM international conference on Multimedia.

Kragel, P. A., Reddan, M. C., LaBar, K. S., & Wager, T. D. (2019). Emotion schemas are embedded in the human visual system. *Science Advances*, 5(7), eaaw4358. doi:10.1126/sciadv.aaw4358

Kramer, A. D., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 201320040.

Lei, S., & Gratch, J. (2019, 3-6 Sept. 2019). *Smiles Signal Surprise in a Social Dilemma*. Paper presented at the 2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII).

Lerner, J. S., & Tiedens, L. Z. (2006). Portrait of the angry decision maker: How appraisal tendencies shape anger's influence on cognition. *Journal of Behavioral Decision Making*, 19, 115-137.

Loewenstein, G., & Lerner, J. S. (2003). The role of affect in decision making. In R. Davidson, H. Goldsmith, & K. Scherer (Eds.), *Handbook of Affective Science* (pp. 619-642): Oxford: Oxford University Press.

Lucas, G., Gratch, J., King, A., & Morency, L.-P. (2014). It's only a computer: Virtual humans increase willingness to disclose. *Computers in Human Behavior*, 37, 94-100.

Mao, W., & Gratch, J. (2012). Modeling Social Causality and Responsibility Judgment in Multi-Agent Interactions. *Journal of Artificial Intelligence Research*, 44, 223-273.

Marsella, S., & Gratch, J. (2003). *Modeling coping behaviors in virtual humans: Don't worry, be happy*. Paper presented at the Second International Joint Conference on Autonomous Agents and Multi-agent Systems, Melbourne, Australia.

Marsella, S., & Gratch, J. (2009). EMA: A process model of appraisal dynamics. *Journal of Cognitive Systems Research, 10*(1), 70-90.

Marsella, S., Gratch, J., & Petta, P. (2010). Computational Models of Emotion. In K. R. Scherer, T. Bänziger, & E. Roesch (Eds.), *A blueprint for affective computing: A sourcebook and manual* (pp. 21-46). New York: Oxford University Press.

Mellers, B. A., Schwartz, A., Ho, K., & Ritov, I. (1997). Decision affect theory: How we feel about risky options. *Psychological Science, 8*, 423-429.

Menninghaus, W., Wagner, V., Wassiliwizky, E., Schindler, I., Hanich, J., Jacobsen, T., & Koelsch, S. (2019). What are aesthetic emotions? *Psychological Review, 126*(2), 171.

Moerland, T. M., Broekens, J., & Jonker, C. M. (2016). *Fear and Hope Emerge from Anticipation in Model-Based Reinforcement Learning*. Paper presented at the IJCAI.

Morén, J., & Balkenius, C. (2000). A computational model of emotional learning in the amygdala. *From animals to animats, 6*, 115-124.

Mühl, C., Heylen, D., & Nijholt, A. (2015). Affective brain-computer interfaces: neuroscientific approaches to affect detection *Oxford Handbook of Affective Computing* (pp. 217-232): Oxford University Press Oxford.

Nagamachi, M. (1995). Kansei engineering: a new ergonomic consumer-oriented technology for product development. *International Journal of Industrial Ergonomics, 15*(1), 3-11.

Nardelli, M., Valenza, G., Greco, A., Lanata, A., & Scilingo, E. P. (2015). Recognizing Emotions Induced by Affective Sounds through Heart Rate Variability. *IEEE Transactions on Affective Computing, 6*(4), 385-394. doi:10.1109/TAFFC.2015.2432810

Neal Reilly, W. S. (1996). *Believable Social and Emotional Agents* (CMU-CS-96-138). Retrieved from Pittsburgh, PA:

Neubauer, C., Chollet, M., Mozgai, S., Dennison, M., Khooshabeh, P., & Scherer, S. (2017). *The relationship between task-induced stress, vocal changes, and physiological state during a dyadic team task*. Paper presented at the Proceedings of the 19th ACM International Conference on Multimodal Interaction, Glasgow, UK. <https://doi.org/10.1145/3136755.3136804>

Ochs, M., Niewiadomski, R., & Pelachaud, C. (2015). 18 Facial Expressions of Emotions for Virtual Characters *The Oxford Handbook of Affective Computing* (pp. 261): Oxford University Press, USA.

Ono, D., Honma, K.-i., & Honma, S. (2015). Circadian and ultradian rhythms of clock gene expression in the suprachiasmatic nucleus of freely moving mice. *Scientific Reports, 5*(1), 12310. doi:10.1038/srep12310

Ortony, A., Clore, G., & Collins, A. (1988). *The Cognitive Structure of Emotions*. Melbourne, Australia: Cambridge University Press.

Osgood, C. E., May, W. H., Miron, M. S., & Miron, M. S. (1975). *Cross-cultural universals of affective meaning* (Vol. 1): University of Illinois Press.

Parkinson, B., Fischer, A. H., & Manstead, A. S. R. (2005). *Emotion in social relations: Cultural, group, and interpersonal processes*. New York: Psychology Press.

Partala, T., & Surakka, V. (2003). Pupil size variation as an indication of affective processing. *International Journal of Human-Computer Studies, 59*(1), 185-198. doi:[https://doi.org/10.1016/S1071-5819\(03\)00017-X](https://doi.org/10.1016/S1071-5819(03)00017-X)

Rahman, M. A., Milasi, R. M., Lucas, C., Araabi, B. N., & Radwan, T. S. (2008). Implementation of Emotional Controller for Interior Permanent-Magnet Synchronous Motor Drive. *IEEE Transactions on Industry Applications, 44*(5), 1466-1476. doi:10.1109/TIA.2008.2002206

Reynolds, C., & Picard, R. (2004). *Affective sensors, privacy, and ethical contracts*. Paper presented at the CHI'04 extended abstracts on Human factors in computing systems.

Robinson, D. T., Smith-Lovin, L., & Wisecup, A. K. (2006). Affect control theory *Handbook of the sociology of emotions* (pp. 179-202): Springer.

Rose, J. (2019, Aug 13, 2019). Amazon Says The Face Recognition Tech It Sells to Cops Can Now Detect 'Fear'. *Vice*.

Russell, J. A. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6), 1161-1178.

Sano, A., & Picard, R. W. (2013, 2-5 Sept. 2013). *Stress Recognition Using Wearable Sensors and Mobile Phones*. Paper presented at the 2013 Humaine Association Conference on Affective Computing and Intelligent Interaction.

Scherer, K. R. (2005). Appraisal theory. *Handbook of cognition and emotion*, 637-663.

Scherer, K. R. (2009). Emotions are emergent processes: they require a dynamic computational architecture. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1535), 3459-3474. doi:doi:10.1098/rstb.2009.0141

Scherer, K. R. (2010). Emotion and emotional competence: conceptual and theoretical issues for modelling agents. In K. R. Scherer, T. Bänziger, & E. Roesch (Eds.), *A blueprint for affective computing: A sourcebook and manual*. New York: Oxford University Press.

Shiota, M. N., Neufeld, S. L., Danvers, A. F., Osborne, E. A., Sng, O., & Yee, C. I. (2014). Positive Emotion Differentiation: A Functional Approach. *Social and Personality Psychology Compass*, 8(3), 104-117. doi:10.1111/spc3.12092

Shore, D. M., Rychlowska, M., van der Schalk, J., Parkinson, B., & Manstead, A. S. R. (2018). Intergroup emotional exchange: Ingroup guilt and outgroup anger increase resource allocation in trust games. *Emotion (Washington, D.C.)*. doi:10.1037/emo0000463

Soleymani, M., Aljanaki, A., Wiering, F., & Veltkamp, R. C. (2015, 29 June-3 July 2015). *Content-based music recommendation using underlying music preference structure*. Paper presented at the 2015 IEEE International Conference on Multimedia and Expo (ICME).

Stratou, G., Ghosh, A., Debevec, P., & Morency, L.-P. (2012). Exploring the effect of illumination on automatic expression recognition using the ICT-3DRFE database. *Image and Vision Computing*, 30(10), 728-737. doi:<https://doi.org/10.1016/j.imavis.2012.02.001>

Traum, D., Swartout, W., Marsella, S., & Gratch, J. (2005). *Fight, Flight, or Negotiate*. Paper presented at the Intelligent Virtual Agents, Kos, Greece.

Tsai, J., Bowring, E., Marsella, S., & Tambe, M. (2013). Empirical evaluation of computational fear contagion models in crowd dispersions. *Autonomous Agents and Multi-Agent Systems*, 1-18.

Utz, V., & DiPaola, S. (2020). Using an AI creativity system to explore how aesthetic experiences are processed along the brain's perceptual neural pathways. *Cognitive Systems Research*, 59, 63-72. doi:<https://doi.org/10.1016/j.cogsys.2019.09.012>

van Kleef, G. A., de Dreu, C., & Manstead, A. S. R. (2004). The interpersonal effects of anger and happiness in negotiations. *Journal of Personality and Social Psychology*, 86(1), 57-76.

Wang, K., & Wan, X. (2018). *SentiGAN: Generating Sentimental Texts via Mixture Adversarial Networks*. Paper presented at the IJCAI.

Yechiam, E., Busemeyer, J. R., Stout, J. C., & Bechara, A. (2005). Using Cognitive Models to Map Relations Between Neuropsychological Disorders and Human Decision-Making Deficits. *Psychological Science*, 16(12), 973-978. doi:10.1111/j.1467-9280.2005.01646.x

Zhang, G., Lu, D., & Liu, H. (2018). Strategies to Utilize the Positive Emotional Contagion Optimally in Crowd Evacuation. *IEEE Transactions on Affective Computing*, 1-1. doi:10.1109/TAFFC.2018.2836462