The Ethics of AI and Emotional Intelligence

Data sources, applications, and questions for evaluating ethics risk

Gretchen Greene, CPhil, MS, JD

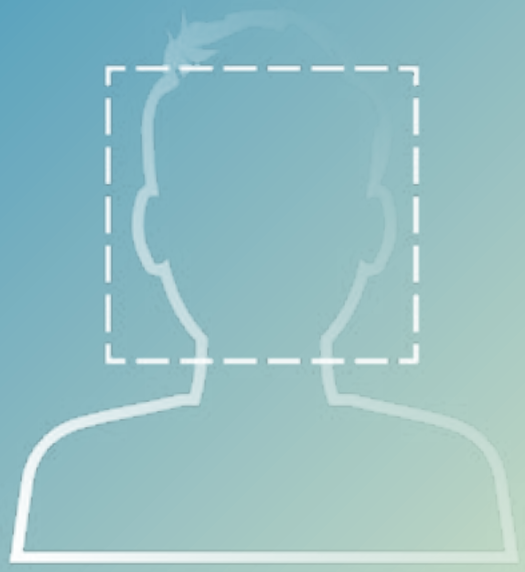

PARTNERSHIP ON AI

Contents

Executive Summary

2019 seemed to mark a turning point in the deployment and public awareness of artificial intelligence designed to recognize emotions and expressions of emotion. The experimental use of AI spread across sectors and moved beyond the internet into the physical world. Stores used AI perceptions of shoppers' moods and interest to display personalized public ads. Schools used AI to quantify student joy and engagement in the classroom. Employers used AI to evaluate job applicants' moods and emotional reactions in automated video interviews and to monitor employees' facial expressions in customer service positions.

It was a year notable for increasing criticism and governance of AI related to emotion and affect. A widely cited review of the literature by Barrett and colleagues questioned the underlying science for the universality of facial expressions and concluded there are insurmountable difficulties in inferring specific emotions reliably from pictures of faces.[1] The affective computing conference ACII added its first panel on the misuses of the technology with the aim of increasing discussions within the technical community on how to improve how their research was impacting society.[2] Surveys on public attitudes in the U.S.[3] and the U.K.[4] found that almost all of those polled found some current advertising and hiring uses of mood detection unacceptable. Some U.S. cities and states started to regulate private[5] and government[6] use of AI related to affect and emotions, including restrictions on them in some data protection legislation and face recognition moratoria. For example, the California Consumer Privacy Act (CCPA), which went into effect January 1, 2020, gives Californians the right to notification about what kinds of data a business is collecting about them and how it is being used and the right to demand that businesses delete their biometric information.[7] Biometric information, as defined in the CCPA, includes many kinds of data that are used to make inferences about emotion or affective state, including imagery of the iris, retina, and face, voice recordings, and keystroke and gait patterns and rhythms.[8]

1 Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Corrigendum: Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. Psychological Science in the Public Interest, 20(3), 165–166. https://doi.org/10.1177/1529100619889954

2 Valstar, M., Gratch, J., Tao, J., Greene, G., & Picard, P. (2019, September 4). Affective computing and the misuse of "our" technology/ science [Panel]. 8th International Conference on Affective Computing & Intelligent Interaction, Cambridge, United Kingdom.

3 Only 15% of Americans polled said it was acceptable for advertisers to use facial recognition technology to see how people respond to public advertising displays. It is unclear whether the 54% of respondents who said it was not acceptable were objecting to the use of facial analysis to detect emotional reaction to ads or the association of identification of an individual through facial recognition with some method of detecting emotional response. See Smith, A. (2019, September 5). More than half of U.S. adults trust law enforcement to use facial recognition responsibly. Pew Research Center. https://www.pewresearch.org/internet/2019/09/05/more-than-half-of-u-s-adults-trust-law-enforcement-to-use-facial-recognition-responsibly/

4 Only 4% of those polled in the U.K. approved of analysing faces (using "facial recognition technologies", which the report defined as including detecting affect) to monitor personality traits and mood of candidates when hiring. Ada Lovelace Institute (2019, September). Beyond face value: public attitudes to facial recognition technology [Report], 11. Retrieved from https://www.adalovelaceinstitute.org/wp-content/uploads/2019/09/Public-attitudes-to-facial-recognition-technology_v.FINAL_.pdf

5 SB-1121 California Consumer Privacy Act of 2018, AB-375 (2018). https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB1121 See also proposed housing bills, No Biometrics Barriers to Housing Act. https://drive.google.com/file/d/1w4ee-poGkDJUkcEMTEAVqHNunplvR087/view [proposed U.S. federal] and Senate bill S5687 [proposed New York state] https://legislation.nysenate.gov/pdf/bills/2019/S5687

6 See Bill S.1385 [MA face recognition bill in process, as of June 23, 2020]. https://malegislature.gov/Bills/191/S1385/Bills/Joint and AB-1215 Body Camera Accountability Act [Bill enacted in CA] https://leginfo.legislature.ca.gov/faces/billCompareClient.xhtml?bill_id=201920200AB1215.

7 The CCPA gives rights to California residents against a corporation or other legal entity operating for the financial benefit of its owners doing business in California that meets a certain revenue or data volume threshold. SB-1121 California Consumer Privacy Act of 2018, AB-375 (2018). https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB1121

8 California Consumer Privacy Act of 2018, AB-375 (2018).

All of this is happening against a backdrop of increasing global discussions, reports, principles, white papers, and government action on responsible, ethical, and trustworthy AI. The OECD's AI Principles, adopted in May 2019 and supported by more than 40 countries, aimed to ensure AI systems would be designed to be robust, safe, fair and trustworthy.[9] In February, 2020, the European Commission released a white paper, "On Artificial Intelligence - A European approach to excellence and trust", setting out policy options for the twin objectives of promoting the uptake of AI and addressing the risks associated with certain uses of AI.[10] In June 2020, the G7 nations and eight other countries launched the Global Partnership on AI, a coalition aimed at ensuring that artificial intelligence is used responsibly, and respects human rights and democratic values.[11]

At its best, if artificial intelligence is able to help individuals better understand and control their own emotional and affective states, including fear, happiness, loneliness, anger, interest and alertness, there is enormous potential for good. It could greatly improve quality of life and help individuals meet long term goals. It could save many lives now lost to suicide, homicide, disease, and accident. It might help us get through the global pandemic and economic crisis.

At its worst, if artificial intelligence can automate the ability to read or control others' emotions, it has substantial implications for economic and political power and individuals' rights.

Governments are thinking hard about AI strategy, policy, and ethics. Now is the time for a broader public debate about the ethics of artificial intelligence and emotional intelligence, while those policies are being written, and while the use of AI for emotions and affect is not yet well entrenched in society. Applications are broad, across many sectors, but most are still in early stages of use.

9 Forty-two countries adopt new OECD Principles on Artificial Intelligence. OECD. Retrieved March 22, 2019, from https://www.oecd.org/science/forty-two-countries-adopt-new-oecd-principles-on-artificial-intelligence.html

10 European Commission. White paper On artificial intelligence - A European approach to excellence and trust, 1. https://ec.europa.eu/info/sites/info/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf

11 Joint statement from founding members of the global partnership on artificial intelligence. Government of Canada. Retrieved July 23, 2020, from https://www.canada.ca/en/innovation-science-economic-development/news/2020/06/joint-statement-from-founding-members-of-the-global-partnership-on-artificial-intelligence.html

Introduction

This report is about artificial intelligence that targets emotions or other affective states and the ethical issues that it raises. It is for anyone who is thinking about how, as a society, we want to develop and use AI, what data we should collect, what inferences we should make, and what safeguards we should put in place. Artificial intelligence designed to recognize, influence, and simulate emotion, interest, and engagement has been increasingly in the marketplace and in the news. Policymakers around the world are considering how to protect privacy, reduce bias, protect workers' and citizens' rights, and ensure that the development and deployment of artificial intelligence is done in a responsible way that benefits society. Individual scientists and companies are thinking about what best practices should be, how their work might be misused,[12] and whether there are applications that should be avoided. Human rights advocates and ethicists are joining the conversation, asking how artificial intelligence is affecting and will affect society.

These discussions have been happening more generally about artificial intelligence, ethics, and policy, and they have also been happening specifically about artificial intelligence related to emotion and affect.[13] This report was informed and inspired by conversations the author had with more than 200 individuals from industry, academia, government, civil society and news media organizations, while leading the work at Partnership on AI on affective computing - computing related to emotions or other affective states[14] - and ethics, over twelve months in 2019 and 2020.[15]

It became clear in the course of those conversations and the affective computing and ethics convenings led by the author in the U.S. and the U.K. in 2019, that a resource was needed that would provide a common starting point. Participants with diverse backgrounds needed a shared understanding of how affective computing is being used, to be able to think about how it should be used. This report creates a foundation for multi-stakeholder conversations, policy making, and public debate about the development, use, benefits and risks of affective computing. Specifically, it presents:

- Definition and categories of affective computing
- Types of human expressions, sensors, and data types used in affective computing
- Categories and examples of current applications of affective computing
- A collection of questions to use as tools for developing best practices for the ethical and beneficial development and deployment of affective computing
- Sources for further reading: research, reports, news, polls, and legislation

12 For example, starting in 2020, a leading machine learning research conference, NeurIPS, is requiring societal impact statements to accompany technical paper submissions. Johnson, K. (2020, February). Neurips requires AI researchers to account for societal impact and financial conflicts of interest. *VentureBeat*. https://venturebeat.com/2020/02/24/neurips-requires-ai-researchers-to-account-for-societal-impact-and-financial-conflicts-of-interest/

13 Smaller discussions of affective computing and ethics started years ago, with the founding of the modern field of affective computing. See Picard, R. W. (1997). Chapter 4: Potential concerns, *Affective computing*. The MIT Press.

14 "Affective computing is computing that relates to, arises from, or deliberately influences emotion or other affective phenomena." Picard, R. W. (1997). *Affective computing*. The MIT Press.

15 A limitation of the research was geographic diversity. Most of the conversations were in the U.S. or the U.K. and all were conducted in English. The author led convenings in the U.S. and the U.K. and had many smaller discussions in both countries. Conversations, including those at the ACII conference, the Harvard Berkman Klein Center, and the Canadian consulate in NYC and in the international human rights community, did include some participants beyond the U.K. and the U.S. but planned meetings in Hong Kong and Brazil were cancelled by COVID-19.

Affective Computing Definition and Language Problems

Affective computing and AI

AI related to emotion sits within the technical field of affective computing.

Affective computing, as defined in Picard's 1997 book that founded the modern field, is "computing that relates to, arises from, or deliberately influences emotion or other affective phenomena."[16] Readers will have a fairly good idea of the scope of the field if they think of it as including any computer program that has something to do with emotions like joy, sadness, fear, and anger or with related affective states like interest, alertness, or engagement. This report focuses on artificial intelligence, but strictly speaking, affective computing does not have to be AI. For example, the definition of affective computing is broad enough to include the tallying of results from an online survey about respondents' mood.[17]

Affective computing includes sensing, recognition, influence, and simulation. AI is being developed to try to do all of these things:

- Sense signals and process data related to emotion
- Find patterns and correlations between emotion and heart rate, perspiration, patterns of speech, direction of gaze, body language, location, pupil dilation, frequency of messages sent and quickness of response.
- Recognize outward expressions associated with emotion, like a smile
- Recognize the emotion a person intends to convey or accidentally conveys
- Replicate an outside observer's perception and predictions about someone else's emotional state, using everything the observer knows, including context
- Recognize or influence the emotion a person feels (both short term states, like happiness or sadness, and longer term states, like depression or anxiety)
- Recognize or influence other affective states including interest, engagement, or alertness
- Measure levels of arousal and negativity/positivity[18] rather than describing a specific emotion
- Simulate emotion or expressions of emotion, e.g. a social robot with gestures meant to convey situationally appropriate emotional responses

16 Picard, R. W. (1997). *Affective computing*. The MIT Press.

17 While AI is sometimes used as a very broad catch-all term, especially for marketing purposes, computer scientists define it more narrowly. In one of the most widely distributed books about AI, *Artificial Intelligence: A Modern Approach*, the authors defined AI as "the designing and building of intelligent agents that receive percepts from the environment and take actions that affect that environment." Russell, S., & Norvig, P. (2013). *Artificial intelligence*. Pearson Education UK.

18 The technical terms from psychology related to this are arousal and valence. See American Psychological Association. *APA Dictionary of Psychology*. Retrieved July 23, 2020, from https://dictionary.apa.org/emotional-valence and https://dictionary.apa.org/arousal

Language problems in public discussions

This report largely avoids using many of the terms which came up in broader discussions of affective computing and which are more familiar to some, terms like emotion recognition, affect recognition and emotion AI. It became apparent early on in the author's research that which terms are used and how they are used or defined are inconsistent across users and sometimes even within a single report.[19] This causes significant confusion about what is being talked about. It makes criticisms seem broader than they are. It makes face recognition and AI related to emotions seem connected in ways that they are not. It is a problem that should be addressed. However, finding consensus on what language should be used proved to be very difficult.

As an example of one of the difficulties language creates, consider the phrase "emotion recognition." It seems at first glance to have an obvious meaning in English, but until we agree on a definition, it contains ambiguity that can mask or cause serious problems in communication and understanding.

Interpreted narrowly, many readers might reasonably think that AI doing "emotion recognition" is reading their innermost feelings. Maybe it does or maybe it will in the future, but applications today are far more likely to be reading something much more superficial.

If we consider the example of analyzing photographs of faces, the computer program might only be an automated smile detector or it might be trying to predict how a person looking at a photograph would have said the person in the photograph was feeling. Both are very different from being able to predict how the person actually was feeling.

Supervised deep learning models, the kind of computer program often used for analyzing images, use labeled examples to learn the patterns it uses to label future unlabeled examples. The labels for the example photographs for the training set are almost always supplied not by the people in the photographs, but by workers hired to label the photographs.[20] The workers look at each picture of a face and choose from a short list of pre-selected emotion labels, e.g. "happy", "sad", and "angry". The resulting computer program then uses the patterns it has found in the labeled examples to label new photographs. The resulting labels are sometimes called "perceived emotion" rather than "emotion."

A broad definition of emotion recognition might cover detecting a smile, correctly predicting what emotion label a person looking at a photograph of a face would have chosen, or correctly predicting how the person in the photograph was actually feeling. A narrow definition might choose any one of them. In the inconsistent use the author has seen, it is often unclear what is meant, and worse, discussants may not realize they have different definitions in mind.

19 For example, the 2019 *AI Now Report*, as accessed on December 24, 2019 (a different version than was available online on June 23, 2020) contained internally inconsistent definitions of affect recognition. p. 6 "Affect recognition - a subclass of facial recognition that claims to detect things such as personality, emotions, mental health, and other interior states", p. 50 "Affect recognition is an AI-driven technology that claims to be able to detect an individual's emotional state based on the use of computer-vision algorithms to analyze their facial microexpressions, tone of voice, or even their gait" Crawford, K., Dobbe, R., Dryer, T., Fried, G., Green, B., Kaziunas, E., Kak, A., Mathur, V., McElroy, E., Sanchez, A. N., Raji, D., Rankin, J. L., Richardson, R., Schultz, J., West, S. M., & Whittaker, M. (2019). AI Now 2019 report. [Report]. The AI Now Institute.

20 Rarely with images, but with other kinds of input data, like heart rate, researchers often use mood questionnaires filled out by the individuals being monitored to create the training set. In that case, the models are trained from data about how individuals were actually feeling, or at least how they said they were feeling.

Sensors, inputs, inferences and applications

Breadth of signals in use

Many applications that are appearing on the market and in the news use face analysis, which may include facial expression or emotion labeling.[21] Cloud computing services from Google, Microsoft, and Amazon offer or have offered emotion labeling as part of their face analysis functionality.[22]

It is not hard to understand the basic idea of trying to guess someone's mood from their facial expression, even without understanding anything about how computer vision works. It is not hard to understand the difference between a smile and happiness and to see that either one can exist without the other. It makes face analysis a useful example for discussion.

But images of faces are just one kind of data. It is important to understand the breadth[23] and combinations[24] of other kinds of data being used in affective computing. When we think about extracting information about emotion, some kinds of human expressions and behaviors, like voice, are nearly as easy as facial expressions for us to understand.

It may be difficult to describe what it is in a voice that tells us in the first word that a friend is upset, distracted or thrilled to talk to us. But we know. We know on the phone where almost all we have is voice.[25] So it seems plausible that there is something we hear in it: the tone, the speed or the loudness, that is the thing (or combination of things) that we use to sense our friend's emotion, that maybe a computer program can also use.

We have less intuition about some of the other signals computers are using like WiFi, patterns of keyboard use, infrared sensors or electrodermal activity. But our lack of intuition does not mean there are not useful patterns to be found. In fact, they might be even more useful, not just mimicking humans, but providing a kind of information we did not have before.

21 Note that face analysis is not the same as facial recognition. Face analysis is a broad term for AI used to analyze images of faces. It includes face recognition or facial recognition, which tries to identify if two face images are the same person. It also includes other kinds of analysis that use images of the face to try to infer various characteristics or states like gender; age; race; mood; alertness; and disease. See Partnership on AI. (2020, February). Bringing Facial Recognition Systems to Light. [Report] https://www.partnershiponai.org/facial-recognition-systems/ for more information on facial recognition systems.

22 Google. *Detect faces*. Google cloud. Retrieved June 30, 2020, from https://cloud.google.com/vision/docs/detecting-faces. Amazon. *Guidelines on Face Attributes.* AWS. Retrieved June 1, 2020, from https://docs.aws.amazon.com/rekognition/latest/dg/guidance-face-attributes.html. Microsoft. *What is the Azure Face service?* Microsoft. Retrieved June 1, 2020, from https://docs.microsoft.com/en-us/azure/cognitive-services/face/overview.

23 For a review of methods and applications see Calvo, R., & D'Mello, S. (2010). Affect detection: An interdisciplinary review of models, methods, and their applications. *IEEE Transactions on Affective Computing*, 1(1), 18-37.

24 See Sharma, R., Pavlovic, V., & Huang, T. (1998). Toward multimodal human-computer interface. Proceedings of the IEEE, 86(5), 853-869 and Pantic, M., & Rothkrantz, L. (2003). Toward an affect-sensitive multimodal human-computer interaction. *Proceedings of the IEEE*, 91(9), 1370-1390.

25 Research suggests we are better at reading emotion from voice alone than from facial expressions alone or from facial expressions and voice. Seppälä, E. (February 26, 2018). You're more emotionally intelligent on the phone, really. We're better at reading voices than faces. Psychology Today. https://www.psychologytoday.com/us/blog/feeling-it/201802/youre-more-emotionally-intelligent-the-phone-really

Examples of affective computing input data sources

Human expression	Data type	Potential application settings	Examples of extracted features*
Voice26	Audio	Call center audio; audio channels from video calls or meetings; home or car assistants27; social robots	Speed; rhythm; variation in tone; gaps between words; pitch contours and loudness
Facial expression28	Visual	Autonomous vehicles; social robots; industrial robots; surveillance drones; cameras on a user's or others' devices; interior/exterior wall mounted camera/video; social media	The state of facial muscles; facial action units29; combinations of edges; intersections; patterns and relationships that a model learns directly from the images
Body language and posture	Visual	Video cameras (public/private space including car mounted inward or outward); weight sensors in car or other seats	Eye movement; openness and pupil dilation; gaze direction; head position; gait; posture; gestures; path of travel
Other physiology	Visual; infrared; wireless signals (radar; WiFi); heart rate; weight sensor	Wearable devices; car seats; medical equipment; video cameras; EEG; EKG	Blood flow; heart rate and heart rate variability; electrodermal activity; breathing patterns; temperature30
Location	Visual; infrared; WiFi	Retail stores; security	Location; time spent in a location; speed; path; dynamic patterns of movement around certain types of locations

26 Zeng, Z., Pantic, M., Roisman, G. & Huang, T. (2009). A Survey of affect recognition methods: audio, visual, and spontaneous expressions. IEEE Transactions on Pattern Analysis and Machine Intelligence, 31(1), 39-58.

27 Johnston, K. (2019, May 28). Amazon testing emotion recognition gadget. Voicebot.ai. https://voicebot.ai/2019/05/28/amazon-testing-emotion-recognition-gadget/

28 For an early review of affective computing, the proceedings of the 1st International Conference on Affective Computing and Intelligent Interaction (ACII 2005) held in Beijing, China, on 22-24 October 2005 included, among other areas of research, affective face, gesture and speech processing. Tao, J., & Tan, T. (2005). Affective computing: A review. Affective Computing And Intelligent Interaction, Proceedings, 3784, 981-995.

29 The Facial Action Coding System (FACS) breaks down facial expressions into components of muscle movement, called Action Units. See Facial action coding system. In Wikipedia. Retrieved July 23, 2020, from https://en.wikipedia.org/wiki/Facial_Action_Coding_System

30 Saini, T., & Bedekar, M. (2017, June 15). Inferring user emotions using physiological signals from mouse and keyboard. International Conference on Intelligent Computing and Control Systems 2017, Madurai. https://www.researchgate.net/publication/316587752_Inferring_User_Emotions_Using_Physiological_Signals_from_Mouse_and_Keyboard

Human expression	Data type	Potential application settings	Examples of extracted features*
Choice of words31	Text or symbols	Chat logs; messages and documents written by the user; responses written by other users; [ML generated] transcripts of audio conversations	Explicit statements of emotion; tone from choice of words; choice of emoji
User device and internet interaction patternsa		Operating system and UI interaction; browsing history & mouse movement; social media interactions; communication and location patterns	Number and time of messages sent; keyboard force32 and use patterns; amount of time spent scrolling through and scanning content; social media likes and reactions

*Not comprehensive and may be oversimplified. A deep learning computer vision model is often learning combinations of edges, intersections, patterns and relationships directly from the images without providing any explanation or explicit features. A model trained to identify "sleepy" faces, for example, might have picked up on an important pattern around the eyes for how faces labeled "not alert" are different from faces labeled "alert" without anyone ever defining eye or eye openness.

31 See Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. Foundations and Trends® in Information Retrieval, 2(1–2), 1-135 and Strapparava, C., & Mihalcea, R. (2008). Learning to identify emotions in text. Proceedings of the 2008 ACM Symposium on Applied Computing, 1556-1560.

32 Saini, T., & Bedekar, M. (2017, June 15). Inferring user emotions using physiological signals from mouse and keyboard. International Conference on Intelligent Computing and Control Systems 2017, Madurai. https://www.researchgate.net/publication/316587752_Inferring_User_Emotions_Using_Physiological_Signals_from_Mouse_and_Keyboard

Inferences and Applications

The applications of affective computing are as broad as the data sources. There are applications in health, education, entertainment, advertising, employment decisions, law enforcement, transportation and agriculture, ranging from monitoring driver alertness to detecting depression.

Those listed in this section came from interviews with affective computing scientists, engineers, and entrepreneurs and from other conversations. They came from searches of news media, academic journals, and websites and from conference talks. They came from data collection and use notices required by the California Consumer Protection Act and from filed patents. The list is not comprehensive, but it will give the reader a good sense of the breadth of applications that are being considered or piloted or are already in broader use. Some of these applications are very controversial. Others may be seen as beneficial or as important steps towards important future beneficial applications. No value judgment is intended by either the inclusion or exclusion of any application here.

Detecting, treating or assisting with disease or disability

Many mental and physical conditions, including autism, depression, and stroke, impact the ability to understand, control or communicate emotions. AI is being applied to assess and treat a variety of conditions and to improve communication related to emotions.

- Facial analysis, wearable sensors and other data are used to provide emotion training,[33] identify increasing stress and stressors, and improve communication and understanding[34] for people with autism and caregivers.

- Deep learning research shows promise in assessing and treating depression,[35] anxiety[36] and other mental health conditions, using a variety of kinds of data as inputs. The number of papers on applications of AI in depression studies and interventions went up from 6 in 2010 to 117 in 2018.[37]

- Wearables are being developed for stress evaluation and relief. One prototype is a wearable cap that uses resting heart rate and EEG to measure stress and responds with soothing vibration and sounds.[38]

33 Metz, C. (2019, July). Google Glass may have an afterlife as a device to teach autistic children. New York Times. https://www.nytimes.com/2019/07/17/technology/google-glass-device-treat-autism.html

34 Picard, R. (2009). Future affective technology for autism and emotion communication. Philosophical Transactions of the Royal Society B, 364(1535), 3575-3584.

35 S'Adan, M., Pampouchidou, A., & Meriaudeau, F. (2018). Deep learning techniques for depression assessment. 2018 International Conference on Intelligent and Advanced System (ICIAS), 1-5.

36 Pintelas, E., Kotsilieris, T., Livieris, I., & Pintelas, P. (2018). A review of machine learning prediction methods for anxiety disorders. Proceedings of the 8th International Conference on Software Development and Technologies for Enhancing Accessibility and Fighting Info-exclusion, 8-15.

37 Mcintyre, R., Latkin, C., Vu, G., Gwee, K., Ho, C., & Ho, R. (2019). The current research landscape on the artificial intelligence application in the management of depressive disorders: A bibliometric analysis. International Journal of Environmental Research and Public Health, 16(12).

38 Sethi, K., Ramya, T., Singh, H., & Dutta, R. (2019). Stress detection and relief using wearable physiological sensors. TELKOMNIKA, 17(3), 1139-1146.

- Researchers are capturing facial expressions to detect pain, to allow communication when patients cannot communicate pain verbally.[39]

- A reconstructive surgeon uses computer vision facial analysis to evaluate pre and post surgery facial expressions for facial palsy and stroke patients.[40]

- Affective computing has applications to quantified self mood tracking, for individuals who want to gain insights about themselves or, through measuring and tracking their own emotional responses, gain insights into others. One doctor uses a (non AI) comparison of his own mood just before and just after meeting a new patient to guide his conversation about their emotional state.[41]

Additional health applications enabled by affect recognition technologies:

- Facial analysis to diagnose DiGeorge, Down Syndrome and other diseases that affect the formation of the face[42]

- Facial analysis to detect stroke[43] or Parkinson's

- Speech analysis to detect Parkinson's[44], concussions,[45] heart disease, PTSD, TBI[46] concussions[47]

- Controlling wheelchairs using facial gestures[48]

- Describing visual environments including facial expressions for blind users[49]

- Seizure detection and alert systems

- Migraine, MS, addiction, schizophrenia

- Machine learning to automate the detection of movement compensations in stroke patients for patient feedback during rehabilitation exercises.[50]

39 Menchetti, G., Chen, Z., Wilkie, D., Ansari, R., Yardimci, Y., & Cetin, A. (2019). Pain detection from facial videos using two-stage deep learning. 2019 IEEE Global Conference on Signal and Information Processing (GlobalSIP), 1-5.

40 Dusseldorp, J. R., Guarin, D. L., Van Veen, M. M., Jowett, N. A., & Hadlock, T. (2019). In the eye of the beholder: changes in perceived emotion expression after smile reanimation. Plastic and Reconstructive Surgery, 144(2), 457471.

41 Carmichael, A. (2013, February 14). Exploring the future of mood tracking (get your mood on: final part). Quantified self. https://quantifiedself.com/blog/exploring-the-future-of-mood-tracking-get-your-mood-on-finalpart/

42 Mjoseth, J. (2017, March). Facial recognition software helps diagnose rare genetic disease. National Human Genome Research Institute. https://www.genome.gov/news/news-release/Facial-recognition-software-helpsdiagnose-rare-genetic-disease

43 NeuroNews (2019, March 7). Artificial intelligence in prehospital stroke detection: Automation in motion. https://neuronewsinternational.com/artificial-intelligence-prehospital-stroke/

44 Wu, Y., Chen, P., Yao, Yuchen, Y., Xiaoquan, X., Yugui, L., Lifang, . . . Chen, J. (2017). Dysphonic voice pattern analysis of patients in Parkinson's disease using minimum interclass probability risk feature selection and bagging ensemble learning methods. Computational and Mathematical Methods in Medicine, 11.

45 Sennarr, K. (2019, May 20). Artificial intelligence for dementia diagnosis - genetic analysis, speech analysis, and more. Emerj. https://emerj.com/ai-sector-overviews/artificial-intelligence-dementia-diagnosis-genetic-analysisspeech-analysis/

46 Mullin, E. (2017, January). Voice analysis tech could diagnose disease. MIT Technology Review. https://www.technologyreview.com/s/603200/voice-analysis-tech-could-diagnose-disease/

47 Poellabauer, C., Yadav, N., Daudet, L., Schneider, S. L., Busso, C., Flynn, P. J. (2015). Challenges in concussion detection using vocal acoustic biomarkers. IEEE Access. https://ecs.utdallas.edu/research/researchlabs/msplab/publications/Poellabauer_2015.pdf

48 Baig, E. C. (2018, December). A smile can move this motorized wheelchair. USA Today. https://www.usatoday.com/story/tech/talkingtech/2018/12/03/ai-wheelchair-controlled-smile-other-facialexpressions/2183372002/

49 Accenture. (2017, July 28). Accenture develops artificial intelligence-powered solution to help improve how visually impaired people live and work. Accenture Newsroom. https://newsroom.accenture.com/news/accenturedevelops-artificial-intelligence-powered-solution-to-help-improve-how-visually-impaired-people-live-andwork.htm

50 Kashi, S., Feingold-Polak, R., Lerner, B., Rokach, L., & Levy-Tzedek, S. (2020). A machine-learning model for automatic detection of movement compensations in stroke patients. IEEE Transactions on Emerging Topics in Computing, 1.

Agriculture

Related to medical use for humans, voice and face analysis have been used to detect fear, pain, and disease in livestock.

- Computer vision and facial analysis in sheep helps farmers detect pain and detect disease earlier, to provide treatment sooner and to keep it from spreading.[51]

- The automated detection of piglet squeals alert farmers an animal might be in danger[52]

Social robots

Autonomous robots that interact and communicate with humans by following social behaviors and rules are finding possible applications in education and health care and in filling a need for interaction and companionship. Social robots may be designed to detect, respond to, and influence user affect, and display what would be perceived by the user as emotionally appropriate expressions or reactions. They might serve roles as companions, caregivers, pets, sex aids, home assistants, therapy coaches, tutors, or catalysts for human to human social interactions.

- Researchers countered loneliness by using a friendly social robot as a group conversational catalyst in elder care facilities.[53]

- Japan's Ministry of Economy, Trade and Industry has predicted that the Japanese market for nursing care and disabled aid robots will grow to 400 billion yen ($3.8 billion) by 2035.[54]

- A social robot at home significantly reduced symptoms of depression in a majority of older adults with chronic depression living alone.[55]

- Studies have found PARO, a social robot that is a plush baby seal, helps reduce behavioral and psychological symptoms of dementia.[56] It was classified as a Class 2 medical device by U.S. regulators in 2009, but adoption by care facilities has been slow.[57]

- Children with autism talked more with a human companion while interacting with a social robot dinosaur than while interacting with either a third person or with a computer game.[58]

51 McLennan, K., Mahmoud, M., (2019). Development of an automated pain facial expression detection system for sheep (Ovis Aries). Animals 2019, 9(4), 196. https://doi.org/10.3390/ani9040196

52 Houser, K. (2018, February 17). Chinese farmers are using AI to track and monitor pigs. It's like Animal Farm meets 1984. Futurism. https://futurism.com/at-tracks-pigs-chinese-farmers

53 Ostrowski, A., Dipaola, D., Partridge, E., Park, H., & Breazeal, C. (2019). Older adults living with social robots: promoting social connectedness in long-term communities. IEEE Robotics & Automation Magazine, 26(2), 59-70.

54 Foster, M. (2018, March 7). Aging Japan: Robots may have role in future of elder care. Reuters. https://www.reuters.com/article/us-japan-ageing-robots-widerimage/aging-japan-robots-may-have-role-infuture-of-elder-care-idUSKBN1H33AB

55 Bennett, C., Sabanovic, S., Piatt, J., Nagata, S., Eldridge, L., & Randall, N. (2017). A robot a day keeps the blues away. 2017 IEEE International Conference on Healthcare Informatics (ICHI), 536-540.

56 Hung, L., Liu, C., Woldum, E., Au-Yeung, A., Berndt, A., Wallsworth, C., . . . Chaudhury, H. (2019). The benefits of and barriers to using a social robot PARO in care settings: A scoping review. BMC Geriatrics, 19(1), 232.

57 Tergesen, A., & Inada, M. (2010). It's not a stuffed animal, it's a $6,000 medical device; Paro the Robo-Seal aims to comfort elderly, but is it ethical? [Column]. The Wall Street Journal Eastern Edition, A1.

58 Kim, E., Berkovits, S., Bernier, L., Leyzberg, D., Shic, E., Paul, P., & Scassellati, D. (2013). Social robots as embedded reinforcers of social behavior in children with autism. Journal of Autism and Developmental Disorders, 43(5), 1038-1049.

- Last year, Japan started rolling out robots in 500 classrooms to help teach English after a 250m yen ($2.3m) investment from the Education Ministry.[59]

- Playing puzzle games with a social robot programmed with behaviors suggestive of a growth mindset led children to try harder during a challenging task than the children who played with a robot programmed to display a more neutral mindset.[60]

- Children ascribe mental states, the ability to form social relationships, and moral standing to robots.[61]

- Physical robots yield significantly more compliance to requests, increase task performance, and produce more rapid learning, compared to virtual agents, including the same robots displayed on screens.[62]

Education and audience engagement

Besides social robots, there are other ways AI related to emotion, engagement, and interest is being used in education. AI is used to measure audience engagement and happiness, online and in person. It can be used to evaluate students, teachers, public speakers, or events, or to improve content or timing, either in real time or for the next event. In a very different kind of application, some college admissions departments are using AI to judge prospective student interest in attending the school.

In a very different kind of education industry application, the algorithmic scoring of colleges has driven colleges to use "demonstrated interest" algorithms to try to avoid accepting students who will reject their offers.

- Voice and face analysis are being used to measure student engagement and joy in video tutoring and classrooms in China.[63]

- Video cameras and AI capture audience emotional reactions and engagement at conferences, training events, and trade shows. The technology is used to show event planners, speakers, and exhibitors how successful they were and help them improve in the future. With faster analysis, it could allow for real time interventions to recapture audience attention.[64]

- Colleges are basing admissions decisions in part on AI predictions of students' interest in the college, based on factors including how early in high school students interacted with the college's website and how quickly they responded to interview invitations.[65]

59 Lufkin, B. (2020, February). What the world can learn from Japan's robots. BBC. https://www.bbc.com/worklife/article/20200205-what-the-world-can-learn-from-japans-robots

60 Park, H. W., Rosenberg-Kima, R., Rosenberg, M., Gordon, G., & Breazeal, C. (2017). Growing growth mindset with a social robot peer. 2017 ACM/IEEE International Conference on Human-Robot Interaction, 137-145.

61 Kahn, P. H., Kanda, T., Ishiguro, H., Freier, N. G., Severson, R. L., Gill, B. T., . . . Shen, S. (2012). "Robovie, you'll have to go into the closet now": Children's social and moral relationships with a humanoid robot. Developmental Psychology, 48(2), 303-314.

62 Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B., & Tanaka, F. (2018, August). Social robots for education: A review. Science Robotics, 3(21), eaat5954. https://doi.org/10.1126/scirobotics.aat5954

63 Hao, K., (2019, August). China has started a grand experiment in AI education. It could reshape how the world learns. MIT Technology Review. https://www.technologyreview.com/s/614057/china-squirrel-has-started-a-grandexperiment-in-ai-education-it-could-reshape-how-the/

64 Kinnersley, H. (2019, April 5). Reading the room. MeetingsNet. https://www.meetingsnet.com/eventtech/reading-room

65 Belkin, D. (2019, January). Colleges mine data on their applicants. The Wall Street Journal. https://www.wsj.com/articles/the-data-colleges-collect-on-applicants-11548507602

Gaming, movies and entertainment

Entertainment is a closely related application to education. Movie and video game studios and amusement parks are experimenting with ways to improve user experiences. They are automating measurement of audience emotional reactions, creating game characters which display appropriate emotional reactions, and letting user emotions drive virtual (and maybe physical) entertainment experiences.

- Disney and 20th Century Fox are experimenting with using infrared cameras and wearables to analyze facial expressions,[66] heart rate, skin moisture and body movement,[67] to understand audience emotions during theater test screenings, with applications in editing and targeted marketing.
- Crowd reactions and player's facial expressions were used to automatically identify tennis match highlights during live TV coverage of the 2017 U.S. Open Tennis Championships.[68]
- Game company EA endows its FIFA-15 characters with attitudes towards each of the others and over 600 emotional reactions so that the 22 players on the field can display appropriate facial expressions and body language based on what happens in the match.[69]
- A biofeedback-enhanced adventure thriller game uses a standard webcam and/or heart rate sensors and makes the game harder for users who seem afraid or stressed.[70]
- Disney's patent application for a ride design describes using cameras, biometric sensors, and RFID devices, along with "emotion/attention determination software" to determine a rider's physical path and environment on a ride.[71]

Advertising and retail

Marketing has always been at least in part about trying to understand and appeal to customers' interests and emotions. Companies devote vast resources to trying to maximize user engagement online and brick and mortar retail stores are now using AI to infer customer mood and interest from video, WiFi, motion sensors, to personalize advertising, understand product interest, and improve customer service.

66 Rodriguez, A. (2017, July 26). While you're watching Disney's films at the cinema, Disney can now watch you. Quartz. https://qz.com/1039102/disney-can-now-use-infrared-cameras-to-track-your-reactions-to-films-indarkened-cinemas/

67 Charara, S. (2016, January 18). Hollywood is tracking heart pounding movie scenes with wearable tech. Wareable. https://www.wareable.com/wearable-tech/heart-racing-bear-scenes-the-revenant-2186

68 Rayo, E. A. (2019, November 22). Artificial intelligence at Disney, Viacom, and other entertainment giants. Emerj. https://emerj.com/ai-sector-overviews/ai-at-disney-viacom-and-other-entertainment-giants/

69 FIFA 15 - emotional intelligence. EA. Retrieved July 23, 2020, from https://www.ea.com/games/fifa/news/fifa15-features-emotional-intelligence

70 Nevermindgame. Retrieved July 23, 2020, from https://nevermindgame.com/about

71 Bilbao, R. (2017, January). Disney patent would alter rides immediately based on passenger emotions. Orlando Business Journal. https://www.bizjournals.com/orlando/news/2017/01/30/disney-patent-would-alter-rides-immediatelybased.html

- In 2017, a four story digital billboard in London's Piccadilly Circus chose ads based on the perceived age, gender and mood of people walking by, and the make and model of passing cars in the background.[72]

- In a Chicago Walgreens drug store, smart drink cooler doors used the weather, motion sensors, cameras and AI estimates of customer gender, age, gaze direction and emotional response to determine which ads to display.[73]

- A patent held by Walmart describes estimating a shopper's level of agitation or annoyance from biometric data, including heart rate or blood pressure, extracted from video cameras at the checkout line, and combining it with purchase history information.[74]

- The department store Macy's uses in-store beacons that map nearby Bluetooth-enabled devices, and wi-fi routers to track where and how long shoppers linger, or "dwell time", as a proxy for customer interest and engagement.[75]

- In 2018, 7-11, a global convenience store chain, announced plans to use AI systems that detect customer emotion in its 11,000 stores in Thailand.[76]

Hiring and employment

Affective computing is being used to screen candidates in automated interviews and to monitor and evaluate worker fatigue, happiness, stress, and job performance.

- A video interview analysis company, Hirevue, uses facial expressions to capture moods to evaluate job candidates.[77]

- Employers track workers' moods with software that searches email and Slack for words that may be associated with depression and fatigue, and by asking workers to regularly log their frame of mind with a smiley or frowning face, or track their mood on apps.[78]

72 Moore, M. (2017, October). Personalised ads delivered by the billboard that's got its eye on you. The Times. https://www.thetimes.co.uk/article/personalised-ads-delivered-by-the-billboard-that-s-got-its-eye-on-youhrxwrrb3z

73 O'Reilly, L. (2019, January). Walgreens tests digital cooler doors with cameras to target you with ads. The Wall Street Journal. https://www.wsj.com/articles/walgreens-tests-digital-cooler-doors-with-cameras-to-target-youwith-ads-11547206200 and Stinson, L. (2019, February). Walgreens is testing digital cooler screens that track shopper behavior. Curbed. https://www.curbed.com/2019/2/8/18216172/walgreens-cooler-screens-trackingshopper-behavior

74 United States Patent, US 9,299,084 B2 (2016, March 29). United States Patent and Trademark Office. https://pdfpiw.uspto.gov/.piw?PageNum=0&docid=09299084&IDKey=C8DE7EC2D24B&HomeUrl=http://patft1.us pto.gov/netacgi/nphParser?Sect1=PTO1%26Sect2=HITOFF%26d=PALL%26p=1%26u=%252Fnetahtml%252FPTO%252Fsrchnum.htm%26r=1%26f=G%26l=50%26s1=9299084.PN.%26OS=PN/9299084%26RS=PN/9299084&mod=article_inline

75 Simonite, T. (2020, January). Now stores must tell you how they're tracking your every move. WIRED. https://www.wired.com/story/stores-must-tell-you-how-theyre-tracking/

76 Chan, T. F. (2018, March). 7-Eleven is bringing facial-recognition technology pioneered in China to its 11,000 stores in Thailand. Business Insider. https://www.businessinsider.com/7-eleven-facial-recognition-technologyintroduced-in-thailand-2018-3?r=UK&IR=T

77 Zetlin, M. AI is now analyzing candidates' facial expressions during video job interviews. Inc. Retrieved July 23, 2020, from https://www.inc.com/minda-zetlin/ai-is-now-analyzing-candidates-facial-expressions-during-video-jobinterviews.html

78 Cutter, C., & Feintzeig, R. (2020, March). Smile! Your boss is tracking your happiness. The Wall Street Journal. https://www.wsj.com/articles/smile-your-boss-is-tracking-your-happiness-11583255617

- An early stage happiness monitoring wearable would combine heart-rate data with calendar data to determine which meetings or colleagues at the office cause stress.[79]

- McDonald's, a global fast food chain, is using facial analysis in its Japanese restaurants to detect whether employees are smiling while assisting customers to improve customer service.[80]

Chatbots, call centers, and home/auto voice assistants

Voice and text sentiment analysis informs interactions with call center agents and chatbots. Call centers are using AI to monitor both callers' and agents' mood and tone with interventions ranging from nudges and talking point suggestions to the agent to real time alerts to their supervisors. Chatbots and voice assistants can modulate their tone and responses based on perceived user emotion, can use emotion as a cue to pass the call to a live agent, and can offer emotional relief to some users by getting those users to talk without feeling fearful or judged.

- Call centers are using AI to evaluate whether workers are responding with appropriate empathy to callers.[81] Software suggests talking points for agents, based on perceived customer stress or frustration.[82]

- An auto voice assistant moderates its interaction style in response to the driver's perceived emotional state.[83]

- A home assistant device listens for user frustration and can respond to users in excited and disappointed speaking styles.[84]

- An AI chatbot app plays a role as friend and confidant, with more than 2.5 million downloads.[85]

79 Cutter, C., & Feintzeig, R. (2020, March). Smile! Your boss is tracking your happiness. The Wall Street Journal. https://www.wsj.com/articles/smile-your-boss-is-tracking-your-happiness-11583255617

80 CB Insights. (2019, June). Facial recognition is already here: These are the 30+ US companies testing the technology. CB Insights. https://www.cbinsights.com/research/facial-recognition-technology-us-corporations/

81 Dzieza, J. (2020, February). How hard will the robots make us work? The Verge. https://www.theverge.com/2020/2/27/21155254/automation-robots-unemployment-jobs-vs-human-googleamazon

82 Simonite, T. (2018, March). This call may be monitored for tone and voice. WIRED. https://www.wired.com/story/this-call-may-be-monitored-for-tone-and-emotion/

83 Kinsella, B. (2019, January 13). Nuance Automotive demonstrates driver emotion detection at CES and shows how virtual assistants can become proactive. Voicebot.ai.https://voicebot.ai/2019/01/13/nuance-automotivedemonstrates-driver-emotion-detection-at-ces-and-shows-how-virtual-assistants-can-become-proactive/

84 Schlosser, K. (2019, November). Alexa, why so emotional? Because Amazon has created realistic new voice tones for its assistant. GeekWire. https://www.geekwire.com/2019/alexa-emotional-amazon-created-realistic-newvoice-tones-assistant/

85 Olson, P. (2018, March). This AI has sparked a budding friendship with 2.5 million people. Forbes. https://www.forbes.com/sites/parmyolson/2018/03/08/replika-chatbot-google-machine-learning/#176e70374ffa

Wearables and stress relief

Wearables allow for measurement of some biometric signals, including heart rate and electrodermal activity, which are harder to capture in other ways. Stress detection and relief is a particularly active area for affective computing wearables and apps.

- A wrist-worn wearable and companion app track electrodermal activity, temperature, and movement to predict when a user is anxious, stressed, or about to have an epileptic seizure.[86]
- Many companies are including heart rate analytics in wearables to detect and track stress.[87]
- A wearable device on the wrist analyzes a user's voice for emotion.[88]

Automotive and industrial safety

Drivers are being monitored for fatigue, alertness, and attention, to improve automotive and industrial safety with both traditional and semi-autonomous vehicles.

- Tired and distracted drivers are two leading causes of serious automobile accidents. Automotive companies are using a variety of signals from the car and the driver including steering wheel angle, speed, sensors to detect lane markings, head orientation, and eyelid movements, to evaluate driver alertness and distraction, prompting alerts and suggestions to the driver to get a coffee or take a break.
- General Motors and BMW track a driver's head and eye positions as part of their semiautonomous driving systems, to make sure drivers are watching the road and prepared to take back control when they need to.[89]
- A major oil company, Chevron, is analyzing its truck drivers' faces for signs of fatigue to improve safety and productivity along fuel transport routes.[90]
- Deadman switches have long been used in trains to stop the vehicle if the operator is incapacitated. One low tech system used in China to ensure operators stay alert requires a push of a pedal every 30 seconds or first an alarm and then automatic braking will occur. One study looked at using EEG measurements and machine learning to potentially detect fatigue to be able to intervene sooner.[91]

86 Sawh, M. (2019, May 27). Getting all emotional: Wearables that are trying to monitor how we feel. Wareable. https://www.wareable.com/wearable-tech/wearables-that-track-emotion-7278

87 Sawh, M. (2020, June 2). Stress wearables: best devices that monitor stress and how they work. Wareable. https://www.wareable.com/health-and-wellbeing/stress-monitoring-wearables-explained-7969

88 Johnston, K. (2019, May 28). Amazon testing emotion recognition gadget. Voicebot.ai. https://voicebot.ai/2019/05/28/amazon-testing-emotion-recognition-gadget/

89 Quain, J. (2019, March). Eyes on the road! (Your car is watching). The New York Times. https://www.nytimes.com/2019/03/28/business/autonomous-cars-technology-privacy.html

90 CB Insights. (2019, June). Facial recognition is already here: These are the 30+ US companies testing the technology. CB Insights. https://www.cbinsights.com/research/facial-recognition-technology-us-corporations/

91 Zhang, X., Li, J., Liu, Y., Zhang, Z., Wang, Z., Luo, D., . . . Wang, C. (2017). Design of a fatigue detection system for high-speed trains based on driver vigilance using a wireless wearable EEG. Sensors (Basel, Switzerland), 17(3), 486.

Threat detection/intervention and law enforcement

The automated recognition of emotion and intent has possible applications in law enforcement from identifying threats of violence, terrorism, and suicide, to detecting lies and fraud.

- A social media company uses users' posts and their contacts' responses and alerts local authorities if the company think there is an immediate risk of suicide.[92]
- Researchers and entrepreneurs are creating virtual border agents that use AI to search for deception in patterns of blood flow, subtle movements, and micro-gestures.[93]

Communities, politics, and social networks

From smart cities to online social networks to old fashioned politics, there is an interest in capturing, responding to, and influencing the mood not just of individuals, but of a population or community. AI used to track population mood and to understand patterns of emotional contagion is currently mostly limited to using online data, but it could come to incorporate data captured offline as well.

- An emotion based digital art gallery across Stockholm's metro and train system tries to detect and improve commuters' moods. It uses real-time public data gathered from Google searches, news articles, social media and travel traffic information to capture an overall mood in the population and then displays a picture chosen to relieve whichever negative emotion it detects.[94]
- A social media platform company found that reducing the number of positive (or negative) posts a user would see, led to the user producing fewer positive (or negative) posts.[95]

92 Kaste, M. (2018, November). Facebook increasingly reliant on A.I. to predict suicide risk. npr. https://www.npr.org/2018/11/17/668408122/facebook-increasingly-reliant-on-a-i-to-predict-suicide-risk

93 Kendrick, M. (2019, April). The border guards you can't win over with a smile. BBC. https://www.bbc.com/future/article/20190416-the-ai-border-guards-you-cant-reason-with

94 Uses web searches, social media, traffic, etc. to recognize mood Yalcinkaya, G. (2019, March 25). Digital billboards in Stockholm project positive artworks to cheer up anxious commuters. Dezeen. https://www.dezeen.com/2019/03/25/emotional-art-gallery-stockholm/

95 Kramer, A. D., Guillory, J. E., & Hancock, J. T. (2014, June). Experimental evidence of massive-scale emotional contagion through social networks. Proceedings of the National Academy of Sciences of the United States of America. https://doi.org/10.1073/pnas.1320040111

Questions for discussion

While many of the applications of affective computing are still early in development or deployed only at a small scale, use of the technology has expanded significantly in the last couple of years and has spread across domains. It has reached a point that calls for careful consideration. The technology is being used in schools, in workplaces, on the streets, in stores, in our cars, and in our homes. The companies developing and deploying the technology should be thinking about the consequences, but not only the companies. As a society, we should decide together if, when, and how we want to develop and use AI to sense, recognize, influence, and simulate human emotion and affect.

These questions are meant as a tool - as prompts and catalysts for thought and conversation, for anyone interested in exploring the ethical issues associated with AI and emotional intelligence.

The questions were captured from and inspired by conversations about affective computing and ethics the author had with more than 200 individuals, mostly in the U.S. and the U.K. in 2019 and 2020. Participants included a broad range of stakeholders: affective computing scientists and engineers; data scientists; chief technology officers; founders; journalists; ethicists; psychologists; lawyers; and civil liberties advocates. Organizations represented included large technology companies, start-ups, nonprofits, news media organizations, government agencies, and academic institutions.

The questions are of various kinds. Some suggest a specific requirement or test that could be applied when making decisions about the details of a particular product. Others consider the potential impact of affective computing more generally on society. Some are specific to affective computing. Others are not, and may apply to the impact of other kinds of artificial intelligence or of technology more broadly.

Consensus was not reached in those conversations on how to answer these questions or even that these are necessarily the right questions to be answering. That said, the author believes that this collection includes most of the issues that arose in those conversations and that they can be a useful tool for thinking about what the opportunities and risks are for affective computing and what safeguards should be put in place.

Thinking big

- Is the technology at the right stage of development and deployment for a broad discussion on related ethics issues?
- What are the biggest opportunities and benefits of affective computing for society?
- What are the greatest risks of affective computing for society?
- How is the use of affective computing impacting society?
- Are there affective computing applications that we should completely avoid?

How does affective computing fit in existing frameworks?

- How do affective computing concerns fit within the larger debates about privacy and data collection and employee, customer, user, and citizen monitoring?
- Does affective computing require different processes, ethics considerations, safeguards or restrictions for development, deployment, procurement or use, compared to AI generally or to other products or services?
- What types of laws are impacting affective computing and what kinds of safeguards are they creating?
- What laws or types of laws should govern affective computing data collection, inferences or applications?
- How do current events (COVID-19 pandemic, global economic crisis, global Black Lives Matter demonstrations, U.S. Supreme Court decisions, etc.) affect what we should focus on or how we should make decisions when thinking about affective computing and ethics?

Human vs. machine

- What are the differences in what a human or AI can detect about emotion/affect?
- What should the role of human capability be in judging what is an acceptable use for AI?
- How does endowing robots with humanlike expressions or responses impact human-computer interactions and human-human interactions?

Accuracy and inclusivity

- How good is the science and the technology for affective computing generally and in this case specifically? Is it good enough to use in this application?
- How does whether the accuracy is good or bad affect which ethics issues we should think about?
- How can we evaluate our technology or systems so they work for everyone?
- If it is a new company or early stage product, is it acceptable to build a product that only works for a subset of the population?
- How does affective computing accuracy depend on the kinds of data, populations, predictions and uses?[96]

96 A study found two face analysis products labeled Black faces with more negative emotions compared to white faces. Rhue, L. (2019) Racial influence on automated perceptions of emotions. SSRN Electronic Journal.

- How are emotional reactions, expressions, or social rules about display of emotion in different contexts, different in different cultures and how should that impact how we create or use affective computing?
- What kinds of diversity do the internal teams and external partners have and how might that improve affective computing subgroup accuracy[97] or otherwise lower risk?
- How could measuring interest or mood reveal or negatively impact a status historically linked to and protected from discrimination such as disability, national origin, religion, race, age, sexual orientation, or gender?

Privacy and other rights

- If an individual has or should have the right to decide whether and when to reveal certain information, how does affective computing impact that right?
- Does this use of affective computing impact an important right or opportunity, like access to jobs, housing or education?
- How might affective computing impact human and civil rights to freedom of thought and opinion or freedom against unreasonable search and seizure or self incrimination?[98]
- Does the data being used or the inference being made, e.g. depression detection, reveal sensitive health information or other information that should have special protection?
- Can the data being collected, such as face images or geolocation, be used in ways that threaten other rights like freedom of assembly, speech or religion?
- How can identifying or influencing emotion or affect create risks to privacy even if the software never has access to identifying or identifiable information?
- What data should be held only on a user's device, rather than being stored on central databases or shared with third parties?
- Who has access to the inferences about the data subject's emotions or affect?

97 A data labeling team in Cairo flagged that the data set of faces did not include women wearing a hijab, prompting the company to add that to their data set. El Kaliouby, R., & Colman, C. (2020). Girl Decoded. Crown/Archetype, 306.

98 The U.N. Declaration of Human Rights proclaims freedoms of thought and opinion. Universal Declaration of Human Rights (1948, December). https://www.ohchr.org/en/udhr/documents/udhr_translations/eng.pdf. The Fourth and Fifth Amendments of the U.S. Constitution provide rights against unreasonable search and seizure and to not testify against oneself in a criminal case. Constitution of United States of America Amendment IV (1791). https://www.law.cornell.edu/wex/fourth_amendment. Constitution of United States of America Amendment V (1791). https://www.law.cornell.edu/constitution/fifth_amendment

Autonomy and best interest

- When is "serving the best interest of the data subject" the right standard when designing and deploying affective computing?
- To what extent should data subjects or others affected by affective computing be in control over when and how tech is used to infer or influence their emotions or affective state, and whose responsibility is it to ensure that they are?
- How can we give data subjects agency over inferences about their affect or emotion?
- When could telling data subjects inferences about their own emotion/affect be harmful?

Transparency and communications

- What should communication standards or best practices be for affective computing?
- How does the language we use in affective computing affect beliefs about efficacy and appropriate use?
- How should we communicate the difference between recognizing perceived emotional states (how a third party thinks someone feels) versus internally felt emotional states (how the person themselves would describe their own emotional state)?
- What problems arise if we are not clear about the difference between emotional expressions (smile) versus emotional states (happy)?
- Deep learning models can pick up signals in the data in ways that are not obvious to the developers creating the model. Will the developers themselves know whether or not a machine learning model is implicitly recognizing or responding to affect or emotion?
- How can we build a common language to facilitate discussions of affective computing and ethics across industry, academia, governments, news media, civil liberties organizations, and the general public, to replace the inconsistency in choice of words and meanings that we have now?

Question exploration as a tool for evaluating ethics risk

Below are three examples of using a question from the collection above as a prompt to explore the ethics risks in a proposed application of affective computing or to think across applications about benefits and risks. Different groups will arrive at different answers.

Who has access to the inferences in depression detection and why does it matter?

In depression detection, who has access to the information and how they use it makes all the difference. Detection might be the first step towards self help or medical treatment. It might give someone a name for how they have been feeling which could help them find better resources or feel less isolated. Coupled with other information, a more sophisticated system could use depression detection to help a person identify and change patterns that lead to their depression worsening.

Some of the best musicians, actors, athletes, and writers in the world have talked openly about their struggles with depression,[99] but that has not been enough to remove all reputational and economic risks of disclosure.[100] As long as that is true, depression detection without control of who can access the information, will often be against the interests of the individual.

Besides access, how the data is generated and by whom affects what safeguards currently apply. It would seem that depression information would be protected under laws meant to prevent the disclosure and misuse of medical information, but it might not be. In the U.S., The Health Insurance Portability and Accountability Act of 1996 (HIPAA), for example, only applies to traditional health care actors, but most of the products that are exploring the use of AI for emotion detection are not in health care. A music recommendation system, online game, home assistant, fitness watch or interview/hiring software might be capable of detecting depression, through information it collects about user speech or behavior. If it does, a law that only applies to health plans, health care clearinghouses,[101] certain health care providers, and their business associates, as HIPAA does,[102] will almost certainly not apply.[103]

99 For example, Olympian Michael Phelps, and rock stars Bruce Springsteen and Lady Gaga have all talked publicly about their depression. Bain, L., & Jeon, H. (2020, February). What 22 celebrities have said about having depression, in their own words. Good Housekeeping. https://www.goodhousekeeping.com/health/wellness/g3541/celebrities-with-depression/?slide=1

100 Lasalvia, A., Zoppei, S., Van Bortel, T., Bonetto, C., Cristofalo, D., Wahlbeck, K., . . . Thornicroft, G. (2013). Global pattern of experienced and anticipated discrimination reported by people with major depressive disorder: A crosssectional survey. The Lancet, 381(9860), 55-62.

101 A category defined in HIPAA including billing service and community health management information systems. The Health Insurance Portability and Accountability Act of 1996, Publ. L. No. 104-191, 110 Stat. 1936 (1996).

102 HIPAA Administrative Simplification Regulation Text (unofficial version, as amended through March 26, 2013). U.S. Department of Health and Human Services Office for Civil Rights. https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/combined/hipaa-simplification201303.pdf

103 This statement is not offered as legal advice regarding the applicability of HIPAA to any specific set of facts.

What are the biggest opportunities and benefits of affective computing for society?

If AI is able to help individuals better understand and control their own emotional and affective states, there is enormous potential for good. It could greatly improve quality of life and help individuals meet long term goals. It could save many lives now lost to suicide, homicide, disease, and accident.

Widespread suffering from loneliness has been called an epidemic.[104] Distraction and fatigue cause a significant percentage of motor vehicle crashes and workplace accidents.[105] These are two large scale problems that affective computing may be able to help address.

Improving individuals' abilities to expand and control their own levels of interest and engagement could also have broad benefits, contributing to their learning, achievement, and feelings of fulfillment and satisfaction.[106]

What are the greatest risks of affective computing for society?

Affective computing intersects with several areas that societies and laws have singled out for additional protection in other contexts. Affective computing often uses biometric data. It may reveal information about physical or mental health, thoughts, or feelings an individual has not chosen to share. It may interfere with the formation or development of beliefs, ideas, opinions, and identity, by trying to influence individuals' emotions or interest or by incentivizing heightened efforts by individuals to conceal their feelings or to avoid certain stimuli.

The prospect of automated detection of others' emotions adds to worries about AI's potential for pervasive, remote, and cheap monitoring and tracking at scale. Automated influence should be of even more concern. Emotion is a very powerful motivator, driving action. If AI can reach our internal state of mind, it can be used as a tool for influence and control, driving markets, politics, and violence. Much simpler methods, like low tech fake news, have been responsible for driving mobs to violence and murder around the world.[107]

Most applications of affective computing being piloted and deployed right now for industry use are doing comparatively simple and less worrisome things, more on the surface, like detecting smiles or whether a driver's gaze is on the road. But there are a number of risks to these kinds of applications too. One is that they may claim to be doing more than they are.

104 Hessing, T. The curse of loneliness, a modern epidemic. Blue Zones. Retrieved July 20, 2020, from https://www.bluezones.com/2018/12/the-curse-of-loneliness-a-modern-epidemic/.

105 A Virginia Tech study says fatigue causes 20 percent of car crashes. https://www.ehstoday.com/safety/article/21917988/wake-up-and-drive-fatigue-causes-20-percent-of-crashes Nine percent of fatal crashes in 2016 were reported as distraction-affected. https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812517 Accident and injury rates are 18% greater during evening shifts and 30% greater during night shifts when compared to day shifts. https://www.osha.gov/SLTC/workerfatigue/hazards.html#:~:text=Worker%20fatigue%20increases%20the%20risk, when%20compared%20to%20day%20shifts.&text=Increased%20sleep%20problems%20and%20risk,of%20hours% 20worked%20per%20week.

106 Interest and engagement contribute to learning, achievement, and feelings of fulfillment and satisfaction https://psych.wisc.edu/cmsdocuments/CompassHH.pdf pg. 43.

107 McLaughlin, T. (2018, December). How WhatsApp fuels fake news and violence in India. WIRED. https://www.wired.com/story/how-whatsapp-fuels-fake-news-and-violence-in-india/ and Guy, J. (2019, March 27) Fake news sparks anti-Roma violence in France. CNN. https://www.cnn.com/2019/03/27/europe/paris-fakekidnapping-scli-intl/index.html and Logan, E. B. (2019, December 24). 'Fake news' has killed Nigerians. Can a bill stop the violence? Los Angeles Times

COVID-19, Black Lives Matter protests, and affective computing

There is a particularly timely need for governments, industry, and societies to think about the technologies which most directly relate to the extraordinary current events. Affective computing relates to a number of topical conversations driven by the COVID-19 pandemic and the global Black Lives Matter protests.

AI could alleviate COVID-19 mental health problems

The stress of the pandemic and the economic crisis are expected to create a substantial rise in mental health problems. In addition, ventilator use for COVID-19 treatment is causing hallucinations, confusion, and dissociation and it appears that some COVID-19 patients may be left with permanent health problems and disability, increasing risk for depression. There was already substantial work on using AI to detect and alleviate depression and anxiety before the emergence of the novel coronavirus. If this work is far enough along, then affective computing may be poised to play a role in alleviating the mental health impacts of the pandemic.

Social robots could help reduce human proximity and disease transmission

Efforts to control the spread of disease have driven the U.S. and countries around the world to rapidly restructure society to reduce physical contact and increase distance between people. The high COVID-19 mortality rate for seniors and the difficulty in controlling disease transmission in group living settings may create a strong incentive for accelerating the development and use of social robots for eldercare, to allow seniors to live in their homes longer and to reduce transmission risks in nursing homes.

The shift to remote work and education raises issues of privacy and control and may increase demand for affective computing technology

The enormous shift to remote work, education, and other aspects of life during government shelter in place orders raises issues around privacy and control. It is now possible to monitor, record, and analyze office workers and students through video conferencing and other online interactions that before the pandemic would have been in person or by phone call.[108] This creates more available data for all kinds of AI to use. It also may create more demand for affective computing if, for example, employers, teachers, and event managers worry that they cannot monitor engagement or read their audiences' reactions in the way they did when meetings, lectures, and events were in person.

Pandemic health monitoring and affective computing share the privacy risks of tracking biometric, health, and location data

The pandemic has inspired the rapid development and deployment of remote sensors, phone apps, and AI for quarantine enforcement, contact tracing, and health status screening. There is overlap in the kinds of data being collected and in some of the privacy issues being raised for public health pandemic applications and for affective computing.

The Black Lives Matter protests and affective computing raise overlapping bias issues

The Black Lives Matter protests inspired discussions across the U.S. and the world on racial bias, discrimination, and violence. Research has shown some affective computing applications do not work as well on Black individuals compared to white, bringing up some of the same issues of racial bias raised in the protests.

108 This shift to computers as intermediaries in so many more of our interactions has made a new kind and scale of monitoring possible, even easy, on the technology side. What is possible politically and allowable legally in a specific jurisdiction is a question of national and local politics and law. Based on my knowledge as a data protection and privacy lawyer in the U.S., readers in the U.S. should not generally assume that existing law will forbid all kinds of such monitoring that they would object to.

The Black Lives Matter protests are changing the law around law enforcement camera use and facial analysis

The Black Lives Matter protests prompted U.S. legislators to draft legislation requiring law enforcement use of cameras and restricting law enforcement use of facial recognition technology. A law restricting the use of facial recognition, even if it does not legally restrict facial analysis related to emotion, which some such laws have, may, as a practical matter, result in decreased use, since the capabilities are often bundled.

Conclusion

In the middle of 2020, many things are unclear. 2019 might have been the year of emotion recognition. It might have marked the beginning of widespread deployment. It might lead to a future of remarkable technical achievement. It might only have been the high point for hype. It is too soon to say.

It is not yet clear how effective AI will get at recognizing, influencing, and simulating human emotion and affect. What is clear is that if it improves enough, it will be a very powerful tool, however we choose to wield it. Before we walk further down that path, we should think hard about what it would mean if it worked really well and what the implications could be. We should also think hard about what it would mean if it does not work very well, but we use it anyway.

We should think about which questions will help us the most in exploring the ethical issues and unintended consequences of creating and deploying AI connected to emotion and affect. We should direct our ethical analysis to the applications that are now coming into the marketplace and that are just behind them in the research labs. We should ask how affective computing overlaps with our most pressing needs today. We should look further ahead and more broadly, asking what are the best and worst foreseeable uses of this kind of AI and asking how its widespread use might create societal changes and problems that any single use would not. As a society, we should ask together, if, when, and how we want to develop and use AI to sense, recognize, influence, and simulate human emotion and affect.

Acknowledgments

The author gratefully acknowledges the support of Partnership on AI and the contributions of the many PAI Partners and others who have enthusiastically engaged in thoughtful reflection and conversation on affective computing and ethics, those named here and those unnamed.

Kade Crockford, ACLU MA
Laura Ellis, BBC
Sinead O'Brien, BBC
Brenda Leong, Future of Privacy Forum
Isabel Bolo, The Hastings Center
David Roscoe, The Hastings Center
Dr. Mildred Solomon, The Hastings Center
Dr. Mary Czerwinski, Microsoft Research AI
Dr. Javier Hernandez, Microsoft Research AI
Dr. Emre Kiciman, Microsoft Research AI
Dr. Daniel McDuff, Microsoft Research AI
Dr. Rosalind Picard, MIT Media Lab

Special thanks to Peter Eckersley, Kira Hessekiel, and Katherine Lewis for their contributions to the affective computing project while at Partnership on AI.

The views reflected in this work and any errors or omissions are those of the author alone and do not reflect the views of any of the individual contributors, Partnership on AI, or any other organization.