FACING AFFECT RECOGNITION

SÉBASTIEN KRIER

9 · 18 · 2020

FACING THE RISKS POSED BY AFFECT RECOGNITION TECHNOLOGIES

INTRODUCTION TO AFFECT RECOGNITION

The field of affect recognition has been an area of interest for a long time, dating back to at least 1995, when MIT Media Lab professor Rosalind Picard published the seminal book *Affective Computing*. She argues that giving computers the ability to recognize, understand, and express emotions is necessary to facilitate our interactions with them. Erik Brynjolfsson agrees: **"Machines that can speak that language — the language of emotions — are going to have better, more effective interactions with us."**[1]

FACIAL RECOGNITION USES VECTORS THAT MAP OUT FACIAL CHARACTERISTICS AND COMPARE THE DATA FROM MULTIPLE FACIAL IMAGES TO IDENTIFY A SINGLE PERSON.

AFFECT RECOGNITION TECHNOLOGIES, PROCESS FACIAL EXPRESSIONS AND MUSCLE MOVEMENTS TO PREDICT DIFFERENT EMOTIONAL REACTIONS.

Emotion is also a key component of intelligence, ignored for far too long by researchers focused on cognitive capabilities and reasoning. Minksy famously echoed this sentiment as early as 1986, asking whether machines without emotions could even be described as intelligent.[2]

Affect or emotion recognition technology (used here interchangeably) seeks to predict emotional reactions of individuals by analyzing a number of features such as facial expressions, voice patterns, written text, gesture, posture, eye movements, physiological signals, and neurological immersion levels (e.g., oxytocin measurement). Typically, emotions are detected using machine-learning classifiers and, increasingly, deep-learning models.[3]

Differences between Affect Recognition and Facial Recognition

Concerns from civil society and academia about facial recognition prompted large tech companies to reexamine when and where such tools can be legitimately used. For example, a study by the National Institute of Standards and Technology evaluated 189 software algorithms from 99 developers and found empirical evidence for racial bias in most facial recognition algorithms.[4] These efforts also allowed citizens, judges, and regulators to be better equipped to weigh the benefits and risks of facial recognition technologies. Indeed, in early 2020 the House Committee on Oversight and Reform were presented with a diversity of views and recommendations on ensuring commercial transparency and accuracy for facial recognition

technologies.[5] Affect recognition, on the other hand, is less well understood – and consequently not scrutinized to the same degree as other machine-learning applications.

At a basic level, affect recognition technologies are quite similar to facial recognition tools. Both technologies are typically used on a person's face,[6] while the two technologies can be distinguished as follows:[7]

- Facial recognition uses vectors that map out facial characteristics and compare the data from multiple facial images to identify a single person. Affect recognition technologies, however, process facial expressions and muscle movements to predict different emotional reactions. The context in which a facial expression is displayed also plays an important role in this process.[8]

So instead of matching facial features to an existing database, affect recognition algorithms find correlations between facial features (and, often, other variables) and certain predefined emotions, feelings, and internal states. The Partnership on AI has listed several use cases such as the ability to recognize affective states including level of interest, engagement, and alertness.[9]

Applications

A growing number of applications incorporate some form of emotion recognition. In the United States, the start-up Affectiva is exploring how to increase impulse buys and monitor cars' drivers for drowsiness, discomfort, and even intoxication.[10] Affect recognition is also used by call centers to identify the moods of customers on the phone and adjust how they handle the conversation in real time.[11] IBM is developing a robot equipped with affect recognition capabilities to help mitigate groupthink in future crewed missions in deep space.[12] In both China and the United States, affective computing is used and studied to measure student engagement in video tutoring and classrooms.[13]
The technology is also being used in more contentious ways, for example, by law enforcement to test criminal suspects for signs of deception, by employers to vet job seekers, and by fintech companies to set insurance prices.[14] Some of these will be discussed further below.

Risks and Limitations

The use of affect recognition technologies presents a number of technical limitations and societal risks. **A lot of existing research shows that – at least for the time being – affect recognition models are feeble, biased, and prone to abuse.** Some research finds that emotion recognition is as much an issue of context as it is of recognizing facial expressions.[15] The way people communicate emotions such as anger, happiness, sadness, and fear differs depending on the culture, context, and situation. The same facial expression can also express instances of more than one emotion category.[16]

In a World Economic Forum article, El Kaliouby acknowledges that cultural differences will influence the effectiveness of emotion detection tools.[17] Some cultures smile more, others less – but that does not necessarily mean one is happier than the other. Other attributes, such as gender and class, can also affect how emotions are expressed around the world. These findings highlight the significant risk of inaccuracy and ethical implications for companies operating in a multicultural and globalized market.

In 2019, Timnit Gebru and colleagues at Google Research reviewed a smiling attribute classifier and sought to answer counterfactual questions such as the following: Would the prediction change if this face characteristic had been different? The assumption was that a fair smiling classifier should perform consistently regardless of the gender or gender expression of the individual in the photo. Unsurprisingly, the authors found that this was not always the case.[18]

Humans express emotions through a variety of behaviors that a computer vision model may not always capture. For example, humans modulate their facial muscles, eye gaze, body gestures, gait, and speech tone among other characteristics. And yet according to researchers at Microsoft, a lot of emotion recognition work tends to rely on a single modality.[19] The reason is that multimodal databases in real-world conditions are rare and small, and often limited to a single language (something researchers from Harbin Engineering University in China and UCLA have tried to address).[20] Even when multimodal datasets are used, they are often private and it is unclear how signals are analyzed and balanced.[21]

PwC UK recently rolled out a voluntary program offering an AI wristband to its staff that monitors the wearer's pulse and sleep and exercise patterns, and how they recover from heightened anxiety levels. While collecting stress-level data can be used to help manage employee well-being, it can also be used to make inferences about the productivity of teams who appear less stressed than others. In China, sanitation workers in Nanjing already receive prompts such as "please continue working" from their location-tracking bracelets when they are deemed too inactive.[22]

As elaborated further below, while affect recognition is not yet widely used, governments across the world are increasingly experimenting with the technology. Areas to keep an eye out for are law enforcement, the military, and the judiciary: in 2017, a multimodal study by Wu et al. claimed to have developed an algorithm capable of detecting lies in courtrooms by analyzing a variety of micro-expressions such as "lips protruded" and "head side turn."[23]

THE IMPORTANCE OF RESPONSIBLE DESIGN

A pressing question is what organizations, researchers, and policymakers should do to ensure that affect recognition technologies are developed and used responsibly. The IEEE is one standard-setting organization that has worked on Ethically Aligned Design and seeks to address various issues related to emotions and emotion-like control in interactions between humans and the design of AI systems. This includes analyses and recommendations on how affect varies across human cultures, the particular problems of artifacts designed for caring and private relationships, considerations of how intelligent artifacts may be used for "nudging," how systems can support human flourishing, and appropriate interventions for artifacts designed with inbuilt affective systems.[24]

Researchers emphasize emotion recognition can only estimate how *others* might perceive an individual's emotions or suggest broad, population-based trends (such as one film eliciting, on average, a more positive reaction than another).[25] Users of this technology, particularly within the judiciary and law enforcement, need to be fully aware that such technologies cannot assess an individual's internal emotions and experience.

Christoffer Heckman, assistant professor of computer science at the University of Colorado, Boulder, explains that, ultimately, technology designers and society as a whole need to look carefully at how information from AI systems is injected into decision-making processes[26] – particularly given that humans tend to trust these systems more than other figures of authority.[27] The uncritical reliance on algorithmic outputs is likely to be a key challenge in the years to come. Cressida Dick, the head of London's Metropolitan Police, has previously said police should look to use "augmented intelligence," rather than relying on AI systems entirely.[28] With emotion recognition tools, it is debatable whether they should be used at all.

THE EUROPEAN UNION AND UNITED KINGDOM

The European Commission's white paper on AI does not refer to affect recognition directly, although it is not unlikely that some use cases will likely fall in the "high-risk" side of the proposed regulatory framework (should this regulatory framework ultimately get adopted).[29]

The EU Commission also proposed new safety requirements for vehicles, which include drowsiness detection. It expects that by 2038 these will save more than 25,000 lives and help prevent at least 140,000 severe injuries. This is a positive example of a narrow regulation that could be both impactful and not unduly restrictive. Bosch recently announced new plans to use AI to help drivers avoid distraction, "micro-sleep," and forgetting to buckle up while behind the

wheel.[30] The EU is also funding an affect recognition project at the University of Reading seeking to assess whether difficult-to-recognize emotions have less consistent patterns than emotions that are easier to recognize.[31]

In the United Kingdom, the government has not yet considered emotion recognition, with the exception of an evidence meeting by the All-Party Parliamentary Group on Artificial Intelligence. In 2017, a survey found that more than half (50.6%) of UK citizens are "not OK" with any form of emotion-capture technology, while just under a third (30.6%) are "OK" with it, as long as the emotion-sensitive application does not identify the individual.[35] In August 2020, the *Times* reported that the British police force is exploring retrospective facial recognition and affect recognition tools to spot anger and distress, "although issues of ethics have delayed the Home Office pilot."

THE AFFECT RECOGNITION MARKET

According to VentureBeat, the global emotion detection and recognition market was valued at $12 billion in 2018 and is projected to reach $90 billion by 2024.[32]

Gartner predicts that by 2022, 10% of personal devices will incorporate affect recognition technologies, up from less than 1% in 2018.[33] However, while growing, adoption remains slow: in a *Wired* article, the CEO of Oxygen Forensics explained that their emotion detector capability has so far not proved popular and that demand was limited.[34]

THE UNITED STATES

The United States is leading internationally in terms of market share, fueled partly by the parallel explosion in facial recognition technologies.[36] Leading researchers in the field, such as Picard and El Kaliouby, were frequently trained and have taught at leading US universities.

Encouragingly, concerns about the safety and ethics of emotional recognition are beginning to surface in the private sector. For example, Accenture has recently published a report warning clients that people will have "justifiable concerns" about privacy, emotional manipulation, bias, and data governance.[37]

Amazon's Rekognition system documentation warns that the service "is not a determination of the person's internal emotional state and should not be used in such a way."[38] IBM does not offer emotion recognition as a service and does not appear to be planning to. Google's internal ethics review process decided against facial recognition technology, but its Cloud offering allows analyzing photos for facial expressions denoting fear, happiness, surprise, and anger – something Android users can already experiment with.[39]

In terms of the public sector, there are few public examples of use by public institutions, although some stand out: activists have previously protested Amazon's pitch of its Rekognition system to Immigration and Customs Enforcement.[40] DARPA, the Pentagon's "blue sky research" arm known for its breakthrough research in national security technology, has recently invested $1 million to build a machine-learning tool to decode and predict the emotions of allies and enemies.[41]

In 2007, the Transportation Security Administration introduced a program that trained staff to identify potential terrorists via facial expression and behavior. The program was later deemed to be ineffective and lacking scientific evidence by the US Government Accountability Office.[42] The ACLU concluded that the program fueled racial profiling. It taught officers to identify suspects by looking for indicators such as "appearing not to understand questions" and displaying "exaggerated emotions."[43] Automating problematic use cases such as this one will inevitably lead to similar issues but on a far greater scale, as we are already seeing in other contexts.[44]

As emotion recognition tools proliferate, we should expect equivalent scrutiny of how government agencies use them. A piece in the MIT Technology Review observes that unlike facial recognition, there have not been as many dedicated campaigns and working groups for affective computing, and attempts at regulation have been limited.[45] One exception is the state of Illinois' Artificial Intelligence Video Interview Act, which seeks to provide job candidates a bit more insight into how these AI-based hiring tools operate.[46]

The Partnership on AI writes that some US cities and states started to regulate private and government use of AI related to affect and emotions, although this is principally through general data protection legislation and face recognition moratoria. Similarly, legislation protecting biometric information would cover many kinds of data that are used to make inferences about emotions, including imagery of the iris, retina, and face; voice recordings; and keystroke and gait patterns and rhythms.[47]

In its annual report, the AI Now Institute called for a ban on technology designed to recognize people's emotions in certain cases.[48] They write that affect recognition "should not be allowed to play a role in important decisions about human lives, such as who is interviewed or hired for a job, the price of insurance, patient pain assessments, or student performance in school." Asked by Axios, ACLU senior policy analyst Jay Stanley expressed doubts about the science behind emotion recognition.[49]

The Partnership on AI has published a report on the ethics of emotional recognition concluding that society should think hard about the implications of using affect recognition at scale – particularly if it is not effective or accurate.[50] They also report that surveys on public attitudes in

the United States found that almost all of those polled found some current advertising and hiring uses of mood detection unacceptable.

CHINA

It is difficult to find reliable estimates of the size of the affect recognition market in China, although Allied Market Research estimates that the Asia-Pacific region will exhibit the most growth by 2023.[51] According to the *Financial Times* (FT), emotion recognition was a hot topic at China's largest surveillance tech expo in 2019.[52] As facial recognition companies continue expanding their offerings, the affect recognition market will continue growing as well.
Some experts claim that it is unlikely that the technology will be rolled out on a large scale in the next three to five years; however, while it is true that the technology is still in its infancy, a representative from the facial recognition company Megvii told the FT that emotion recognition technology was being widely developed and used within government, particularly public security bureaus.[53]

Indeed, emotion recognition is being deployed at airports and subway stations to identify criminal suspects: Li Xiaoyu, policing expert and party cadre from Xinjiang, told the FT that "Using video footage, emotion recognition can rapidly identify criminal suspects by analyzing their mental state … to prevent illegal acts including terrorism and smuggling."[54] These systems combine other algorithms such as gait recognition, eye tracking, and crowd analysis to make predictions with very little oversight or public scrutiny.[55]

The Ministry of Science and Technology's "Governance Principles for the New Generation Artificial Intelligence" includes requirements such as "AI developers, users and other related stakeholders should have a high sense of social responsibility and self-discipline, and should strictly abide by laws, regulations, ethical principles, technical standards and social norms." and "Through technology advancement and management improvement, prejudices and discriminations should be eliminated as much as possible in the process of data acquisition, algorithm design, technology development, and product development and application.."[56] Such principles, however, have been frequently criticized for being ineffective and unlikely to influence technological development.

While a number of legislative proposals tackle privacy and personal data, emotion recognition does not appear to feature on lawmakers' agendas. Similarly, self-regulation initiatives, standards, and regulatory proposals cover facial recognition technologies[57] but none targeting affect recognition specifically. Yi Zeng et al., in a study published in 2019, called for the industry to create a responsible future and consider issues like informed consent, data revocation, model retraining, risk evaluation, and ethics by design.[58]

According to a survey of 6,152 people by the Nandu Personal Information Protection Research Centre in China, people are growing wary of facial recognition.[59] It is arguable that attitudes toward affective computing will likely follow a similar path. So far, however, there appear to be few or no initiatives from civil society groups and industry associations on the oversight and abuse of affect computing.

ANALYSIS: UNITED STATES–CHINA

Given the field's immaturity and very recent expansion, it is difficult to find clear and comparable trends in the development and utilization of affect recognition technologies across the United States and China. Both ecosystems lead internationally in the research and commercialization of such technologies, while neither China nor the United States currently has any regulations, guidance, or policies relating specifically to the use of affect recognition technologies. Both countries have varying levels of data protection legislation in place, and some US states are exploring regulation proposals. Nor do there appear to be widely adopted industry standards or accepted oversight mechanisms.

In short, the ecosystems are similar to the extent that they are both nascent, lack any meaningful public oversight, and would benefit from further research and development. However both ecosystems differ in two key respects:

I. **Civil society:** In the United States, nonprofits such as AI Now, the Partnership on AI, and the ACLU have already expressed opposition to the widespread use of emotion recognition, and researchers in the field have frequently highlighted the risk of misuse. While surveys indicate growing opposition in China toward facial recognition systems, attitudes toward the oversight of affective computing remain unclear. So far, there appear to be no public calls for regulation or oversight beyond those applicable to AI in general, although recent studies such as that by Yi Zeng et al. cited above are starting to highlight societal implications and practical advice.

II. **Applications:** Most applications in the United States appear to be commercial and intent on understanding and changing consumer behavior – for example, in gaming, the automotive industry, and media. This equally applies to China; however, the Chinese state plays a very active role in incentivizing the development of affect recognition tools in criminal justice and surveillance, and testing these tools in a variety of controversial contexts. This also reflects a broader ideological difference on the role of the state.

Unlike in the US private sector, Chinese companies have been particularly proactive at the UN's International Telecommunication Union to shape standards around facial recognition technologies. This trend might well continue for other technologies such as affect recognition.

The Trump administration's retrenchment from multilateral institutions means standard-setting activities are unlikely to be reflective of the diverse views offered by the US ecosystem.

As Jack Clark predicts, "The challenge for the second half of the decade will be resisting the seemingly unstoppable creep of surveillance-led authoritarian governance."[60] It remains to be seen where, how, and who will lead such efforts.

COOPERATION RATING (GYR)

As observed above, for the time being, there are little to no governance initiatives, regulations, guidance, or laws overseeing the use of emotion detection (whether by public entities or private organizations). Starting with a blank slate might facilitate the co-creation of governance mechanisms; however, ideological, cultural, and political barriers are likely to be significant hurdles to meaningful cooperation.

Overall, the issue of affect recognition governance represents a **Red Light** in the Green-Yellow-Red (GYR) framework put forward by the Asia Society. There are clear differences in how both countries think about these issues: in particular, China's wide deployment of surveillance tools from Xinjiang camps to high schools will likely not be perceived as acceptable by the US government. While some cooperation is possible in R&D, it is unlikely the same would apply to the global governance or regulation of affect recognition systems.

This divergence is already surfacing with facial recognition technologies. The White House's chief technology officer, Michael Kratsios, told the Associated Press that through the Global Partnership on AI (GPAI), the United States seeks to establish shared democratic principles as a counter to China's record of "twisting technology" in ways that threaten human rights, specifically referring to the fact that "Chinese technology companies are attempting to shape international standards on facial recognition and surveillance at the United Nations."[61] The United States is equally concerned about Chinese state-backed companies selling this technology to other autocratic nations such as Venezuela and the UAE. This seems to indicate that at least on the political level, the same position will be held with regard to affect recognition technologies.

Tensions between China and the United States do not appear to be easing. Trump has recently voiced support for further economic decoupling from China[62] – a prospect few American

companies find appealing.[63] Chinese leaders have equally accused the United States of "bullying" others and "hunting" companies.[64] Given these developments, the chances of cooperation and multilateral initiatives on AI governance more generally appear low.

CONCRETE AREAS FOR COLLABORATION

As Jessica Cussins' report concludes, there are areas such as principles and ethical frameworks where China and the United States can fruitfully collaborate. General collaboration on voluntary AI governance and ethics initiatives by private sector organizations, research institutes, and civil society organizations, therefore, should be encouraged to ensure the risks presented by affect recognition are adequately addressed. This includes co-developing oversight mechanisms and funding further research to ensure the responsible development and use of affective computing.

However, as Cussins also observes, "the translation of AI principles into practices surfaces complications as technical decisions collide with economic, social, and political realities."[65] The use and oversight of affect recognition technologies are specific-use cases where such collaboration is likely to surface important divergences in values and ideologies – particularly in the context of population control, crime prevention, and surveillance.

More research in this area is certainly needed. Heckman believes that remedying some of these risks will require active efforts to augment machine learning techniques to consider context (something DARPA is exploring[66]).[67] El Kaliouby thinks the industry will evolve as it acquires more data.[68] The co-creation of new diverse and multimodal datasets is welcomed, and so are research-based collaborations. It is important to remember, however, that better datasets alone will not address the majority of risks and ethical roadblocks.

CONCLUSION

Affect recognition remains an important but nascent field that warrants further development as well as public scrutiny. The research on affect recognition remains generally limited: there are many unknowns in both use and interpretation, and these hurdles should be acknowledged and understood by both developers and deployers. While academics in the field are usually careful about limitations, a lot of this nuance is lost when commercial applications are advertised to potential clients.

There are important risks of discrimination and bias, perhaps even more so than with pure facial recognition systems. Analyzing a face alone is unlikely to be a helpful proxy to more nuanced emotions, and publicly available multimodal approaches and datasets leave a lot to be desired.

Barrett et al. note that the existing tools are unreliable, lack specificity, and are not yet generalizable.[69]

It is therefore arguable that some form of regulation might be required to limit the risk of misuse, and that more research in the field should be encouraged. The Employee Polygraph Protection Act, for example, prohibits most private employers from using polygraphs – a technology that suffers from similar flaws.[70] "Opt-in" requirements and the use of voluntary frameworks such as model cards should also be encouraged.[71]

To conclude, the **three key takeaways** of this report can be summarized as follows:

- **The affect technology market has been rapidly growing in recent years, in part due to the proliferation of facial recognition technologies.** China and the United States are leading in commercializing use cases in a number of industries. However, the technology remains highly imperfect and its capabilities are often overestimated. The academic literature on effectiveness and security is limited, and a lot of claims made by vendors of affect recognition technologies are unsubstantiated or exaggerated.

- **Whether in the private or public sector, the use of emotional recognition is fraught with ethical, design, and political concerns.** As countries start adopting 5G infrastructure, smart cities will enable a host of AI-powered capabilities: in many countries, this will include surveillance tools that seek to predict and infer people's emotions and inner states. The risk of misuse, bias, abuse, and harm will be particularly high in this context, and it is important for affect recognition to be more widely understood, examined, and debated. Given the limitations of affect recognition tools, it is legitimate for individuals to be concerned about how they can influence high-stakes decisions about their lives.

- **For the time being, governance efforts are embryonic at best.** There have been calls for regulation and bringing attention to the risks posed by affect recognition in the United States; however, policy and legislative proposals remain sparse. The Chinese affect recognition market is supported by a proactive industrial policy incentivizing its use as part of the state's surveillance apparatus – as such it seems equally unlikely that oversight will form part of the national agenda anytime soon, let alone on an international bilateral level.

ENDNOTES

[1] "Emotion AI, Explained." MIT Sloan. https://mitsloan.mit.edu/ideas-made-to-matter/emotion-ai-explained.
[2] Minsky, M. L. *The Society of Mind*. New York: Simon and Schuster (1986).
[3] "Emotion Recognition Using Multi-Modal Data and Machine Learning Techniques: A Tutorial and Review." *Information Fusion*, vol. 59, July 2020, pp. 103–26.
[4] Grother, P., Ngan, M., and Hanaoka, K. "Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects." National Institute of Standards and Technology Interagency Internal Report 8280, December 2019. https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf
[5] "Facial Recognition Technology (Part III): Ensuring Commercial Transparency & Accuracy" United States House Committee on Oversight and Reform, January 15, 2020. https://oversight.house.gov/legislation/hearings/facial-recognition-technology-part-iii-ensuring-commercial-transparency
[6] Yang, D., et al. "An Emotion Recognition Model Based on Facial Recognition in Virtual Learning Environment." *Procedia Computer Science*, vol. 125, January 2018, pp. 2–10. ScienceDirect, doi: 10.1016/j.procs.2017.12.003.
[7] "Facial Coding vs. Facial Recognition – What's the Difference?" EnterpriseAI, April 8, 2020. https://www.enterpriseai.news/2020/04/08/facial-coding-vs-facial-recognition-whats-the-difference/.
[8] Kalehoff, Max. "The Clear Line Between Facial Coding and Facial Recognition." https://blog.realeyesit.com/the-clear-line-between-facial-coding-and-facial-recognition.
[9] "The Ethics of AI and Emotional Intelligence." The Partnership on AI, July 30, 2020. https://www.partnershiponai.org/the-ethics-of-ai-and-emotional-intelligence/.
[10] Ford, M. *Architects of Intelligence: The Truth about AI from the People Building It,* Birmingham, UK: Packt Publishing (2018).
[11] "This Call May Be Monitored for Tone and Emotion." *Wired*. www.wired.com. https://www.wired.com/story/this-call-may-be-monitored-for-tone-and-emotion/.
[12] "Emotion-Sensing Robot Launches to Assist Space Station Astronauts." Reuters, December 5, 2019. www.reuters.com, https://www.reuters.com/article/us-space-exploration-robot-idUSKBN1Y92MV.
[13] Kim, Min Kyu, and Tuba Ketenci. "The Role of Expressed Emotions in Online Discussions." *Journal of Research on Technology in Education*, vol. 52, no. 1, January 2020, pp. 95–112. DOI.org (Crossref), doi: 10.1080/15391523.2019.1697861. See also Synced. "EmotionCues: AI Knows Whether Students Are Paying Attention." Synced, January 16, 2020. https://syncedreview.com/2020/01/16/emotioncues-ai-knows-whether-students-are-paying-attention/.

[14] "Emotion-Detecting Tech Should Be Restricted by Law – AI Now." *BBC News*, December 12, 2019. https://www.bbc.com/news/technology-50761116.
[15] Chen, Zhimin, and David Whitney. "Tracking the Affective State of Unseen Persons." *PNAS Proceedings of the National Academy of Sciences of the United States of America,* vol. 116, no. 15, 2019, pp. 7559–64. APA PsycNET, doi: 10.1073/pnas.1812250116.
[16] Barrett, Lisa Feldman, et al. "Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements." *Psychological Science in the Public Interest*, July 2019. Los Angeles, CA, journals.sagepub.com, doi: 10.1177/1529100619832930.
[17] "Computers Can Now Read Your Emotions. Here's Why That's Not as Scary as It Sounds." World Economic Forum. https://www.weforum.org/agenda/2017/03/computers-can-now-read-your-emotions-here-s-why-that-s-not-as-scary-as-it-sounds/.
[18] Denton, Emily, et al. "Detecting Bias with Generative Counterfactual Face Attribute Augmentation." ArXiv:1906.06439 [Cs, Stat], June 2019. arXiv.org. http://arxiv.org/abs/1906.06439.
[19] Kapoor, Ashish, and Rosalind W. Picard. "Multimodal Affect Recognition in Learning Environments." Proceedings of the 13th Annual ACM International Conference on Multimedia, Association for Computing Machinery, 2005, pp. 677–82. ACM Digital Library. doi: 10.1145/1101149.1101300.
[20] Chen, Jing, et al. "HEU Emotion: A Large-Scale Database for Multi-Modal Emotion Recognition in the Wild." ArXiv:2007.12519 [Cs], July 2020. arXiv.org. http://arxiv.org/abs/2007.12519.
[21] Vincent, James. "AI 'Emotion Recognition' Can't Be Trusted." The Verge, July 25, 2019. https://www.theverge.com/2019/7/25/8929793/emotion-recognition-analysis-ai-machine-learning-facial-expression-review.
[22] Hollister, Sean. "These Chinese Sanitation Workers Have to Wear Location-Tracking Bracelets Now." The Verge, April 6, 2019. https://www.theverge.com/2019/4/6/18298562/these-chinese-sanitation-workers-have-to-wear-location-tracking-bracelets-now.
[23] Wu, Zhe, et al. "Deception Detection in Videos." ArXiv:1712.04415 [Cs], December 2017. arXiv.org. http://arxiv.org/abs/1712.04415.
[24] "Affective Computing," The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, IEEE. https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e_affective_computing.pdf.
[25] "Emotion AI Researchers Say Overblown Claims Give Their Work a Bad Name." *MIT Technology Review*. https://www.technologyreview.com/2020/02/14/844765/ai-emotion-recognition-affective-computing-hirevue-regulation-ethics/.
[26] Heckman, Christoffer. "AI Can Now Read Emotions – Should It?" The Conversation, http://theconversation.com/ai-can-now-read-emotions-should-it-128988.
[27] "New Study: 64% of People Trust a Robot More Than Their Manager." https://www.oracle.com/corporate/pressrelease/robots-at-work-101519.html.
[28] Burgess, Matt. "Police Built an AI to Predict Violent Crime. It Was Seriously Flawed." *Wired UK*, August 2020. www.wired.co.uk. https://www.wired.co.uk/article/police-violence-prediction-ndas.

[29] European Commission. "White Paper on Artificial Intelligence: Public Consultation towards a European Approach for Excellence and Trust." Shaping Europe's Digital Future, July 17, 2020. https://ec.europa.eu/digital-single-market/en/news/white-paper-artificial-intelligence-public-consultation-towards-european-approach-excellence.
[30] "Bosch Announces Plans for AI-Powered Car Safety." BusinessCloud, December 16, 2019. https://www.businesscloud.co.uk/news/bosch-announces-plans-for-ai-powered-car-safety/.
[31] "Emotion Recognition: A Statistical Learning Approach." CORDIS. https://cordis.europa.eu/project/id/887283.
[32] "Realeyes Raises $12.4 Million to Help Brands Detect Emotion Using AI on Facial Expressions." VentureBeat, June 6, 2019. https://venturebeat.com/2019/06/06/realeyes-raises-12-4-million-to-help-brands-detect-emotion-using-ai-on-facial-expressions/.
[33] "Gartner Highlights 10 Uses for AI-Powered Smartphones." https://www.gartner.com/en/newsroom/press-releases/2018-03-20-gartner-highlights-10-uses-for-ai-powered-smartphones.
[34] "Amazon Says It Can Detect Fear on Your Face. You Scared?" *Wired*. www.wired.com. https://www.wired.com/story/amazon-detect-fear-face-you-scared/.
[35] "Tech Firms Want to Detect Your Emotions and Expressions, but People Don't like It." Media@LSE, July 11, 2017. https://blogs.lse.ac.uk/medialse/2017/07/11/tech-firms-want-to-detect-your-emotions-and-expressions-but-people-dont-like-it/.
[36] "Emotion Analytics Market Research Report - Global Forecast 2023" | MRFR. https://www.marketresearchfuture.com/reports/emotion-analytics-market-5330. See also "AMR Forecasts $33.9 Billion Emotion Recognition Market." FindBiometrics, February 18, 2020. https://findbiometrics.com/biometrics-news-amr-forecasts-33-9-billion-emotion-recognition-market-021806/.
[37] "Getting Emotional." Accenture, 2020. https://www.accenture.com/_acnmedia/PDF-114/Accenture-Responsible-Use-Of-Emotional-AI-Final.pdf.
[38] "Emotion" Amazon Rekognition Developer Guide, 2020. https://docs.aws.amazon.com/rekognition/latest/dg/API_Emotion.html.
[39] "Amazon Says It Can Detect Fear on Your Face."
[40] "Amazon Says the Face Recognition Tech It Sells to Cops Can Now Detect 'Fear.'" https://www.vice.com/en_us/article/7x59z9/the-facial-recognition-system-amazon-sells-to-cops-can-now-detect-fear.
[41] Brewster, Thomas. "DARPA Pays $1 Million for an AI App That Can Predict an Enemy's Emotions." *Forbes*. https://www.forbes.com/sites/thomasbrewster/2020/07/15/the-pentagons-1-million-question-can-ai-predict-an-enemys-emotions/.
[42] "'Emotion Detection' AI Is a $20 Billion Industry. New Research Says It Can't Do What It Claims." *Washington Post*, July 31, 2019. https://www.washingtonpost.com/business/2019/07/31/emotion-detection-ai-is-billion-industry-new-research-says-it-cant-do-what-it-claims/.
[43] "ACLU v. TSA." American Civil Liberties Union. https://www.aclu.org/cases/aclu-v-tsa.

[44] Corbett-Davies, Sam, et al. "A Computer Program Used for Bail and Sentencing Decisions Was Labeled Biased against Blacks. It's Actually Not That Clear." *Washington Post*, October 17, 2016. https://www.washingtonpost.com/news/monkey-cage/wp/2016/10/17/can-an-algorithm-be-racist-our-analysis-is-more-cautious-than-propublicas/.
[45] "Emotion AI Researchers Say Overblown Claims Give Their Work a Bad Name."
[46] Heilweil, Rebecca. "Illinois Says You Should Know if AI Is Grading Your Online Job Interviews." Vox, January 1, 2020. https://www.vox.com/recode/2020/1/1/21043000/artificial-intelligence-job-applications-illinios-video-interivew-act.
[47] "The Ethics of AI and Emotional Intelligence." The Partnership on AI, 30 July 2020. https://www.partnershiponai.org/the-ethics-of-ai-and-emotional-intelligence/.
[48]"AI Now 2019 Report", AI Now Institute, 2019. https://ainowinstitute.org/reports.html .
[49] Rummler, Orion. "China Touts Dubious Emotion Recognition Tech to Find Criminals." Axios. https://www.axios.com/china-emotion-recognition-f8af57ee-4a1b-43a3-b0fc-b2582b5d9ec6.html.
[50] "The Ethics of AI and Emotional Intelligence."
[51] "AMR Forecasts $33.9 Billion Emotion Recognition Market." FindBiometrics, February 18, 2020, https://findbiometrics.com/biometrics-news-amr-forecasts-33-9-billion-emotion-recognition-market-021806/.
[52] "Emotion Recognition Is China's New Surveillance Craze." *Financial Times*, November 1, 2019. https://www.ft.com/content/68155560-fbd1-11e9-a354-36acbbb0d9b6.
[53] "Emotion Recognition Is China's New Surveillance Craze."
[54] "Emotion Recognition Is China's New Surveillance Craze."
[55] "Emotion Recognition Is China's New Surveillance Craze."
[56] "Governance Principles for the New Generation Artificial Intelligence--Developing Responsible Artificial Intelligence.'." *China Daily*, *June 17, 2019*. https://www.chinadaily.com.cn/a/201906/17/WS5d07486ba3103dbf14328ab7.html.
[57] "Coming into Focus: China's Facial Recognition Regulations." https://www.csis.org/blogs/trustee-china-hand/coming-focus-chinas-facial-recognition-regulations.
[58] Zeng, Yi, et al. "Responsible Facial Recognition and Beyond." ArXiv:1909.12935 [Cs], September 2019. arXiv.org. http://arxiv.org/abs/1909.12935.
[59] "Nearly 80% of People in China Worry about Facial Recognition Data Leaks." *South China Morning Post*, December 9, 2019. https://www.scmp.com/abacus/news-bites/article/3041300/nearly-80-people-china-worry-about-facial-recognition-data-leaks.
[60] CSET Foretell Blog. "COVID-19 Surveillance Strengthens Authoritarian Governments." https://www.cset-foretell.com/blog/surveillance-creep.
[61] O'Brien, Matt. "US Joins G7 Artificial Intelligence Group to Counter China." *Defense News*, May 29, 2020. https://www.defensenews.com/global/the-americas/2020/05/29/us-joins-g7-artificial-intelligence-group-to-counter-china/.
[62] "Trump Raises Prospect of 'Decoupling' US Economy from China." *Financial Times*, September 7, 2020. https://www.ft.com/content/06047bc5-81dd-4475-8678-4b3181d53877.

[63] "US Companies Defy Trump's Threats about 'Decoupling' from China." *Financial Times*, September 9, 2020. https://www.ft.com/content/8d23d65b-ee20-4449-a615-e3d2a9b672f8.
[64] "China Unveils Global Data Security Initiative, Says Some Countries Bullying Others." Reuters, September 8, 2020. www.reuters.com, https://www.reuters.com/article/us-china-security-data-idUSKBN25Z06L.
[65] Cussins, J. N. "AI Principles in Context." Asia Society, August 20, 2020. https://asiasociety.org/sites/default/files/inline-files/Cussins_Principles_Final.pdf.
[66] «DARPA Perspective on AI. » https://www.darpa.mil/about-us/darpa-perspective-on-ai.
[67] Heckman. "AI Can Now Read Emotions – Should It?"
[68] "'Emotion Detection' AI Is a $20 Billion Industry."
[69] Barrett, "Emotional Expressions Reconsidered"
[70] "Emotion AI Researchers Say Overblown Claims Give Their Work a Bad Name."
[71] Mitchell, Margaret, et al. "Model Cards for Model Reporting." Proceedings of the Conference on Fairness, Accountability, and Transparency - FAT* January 2019, pp. 220–29. arXiv.org. doi: 10.1145/3287560.3287596.

Written by Sébastien Krier
Series Edited and Overseen by Heather Evans and Benjamin Cedric Larsen
Cover Artwork Designed by Kisara Moore and Anya Lassila
Traffic Light Design by Anya Lassila

Thank you to the Asia Society Northern California Team for supporting this publication.