

AI can now read emotions—but should it?

January 8 2020, by Christoffer Heckman



Emotion recognition technology, an outgrowth of facial recognition technology, continues to advance quickly. Credit: [Steve Jurvetson/flickr](#), [CC BY-SA](#)

In its [annual report](#), the AI Now Institute, an interdisciplinary research center studying the societal implications of artificial intelligence, called for a ban on technology designed to recognize people's emotions in certain cases. Specifically, the researchers said [affect recognition technology](#), also called emotion recognition technology, should not be used in decisions that "impact people's lives and access to opportunities," such as hiring decisions or pain assessments, because it is not [sufficiently accurate](#) and can lead to biased decisions.

What is this technology, which is already being used and marketed, and why is it raising concerns?

Outgrowth of facial recognition

Researchers have been actively working on [computer vision algorithms](#) that can determine the emotions and intent of humans, along with making other inferences, for at least a decade. Facial expression analysis [has been around since at least 2003](#). Computers have been able to understand emotion [even longer](#). This latest technology relies on the data-centric techniques known as "[machine learning](#)," algorithms that process data to "learn" how to make decisions, to accomplish even more accurate affect recognition.

The challenge of reading emotions

Researchers are always looking to do new things by building on what has been done before. Emotion recognition is enticing because, somehow, [we as humans can accomplish this relatively well from even an early age](#), and yet capably replicating that human skill using computer vision is still challenging. While it's possible to do some pretty remarkable things with images, such as [stylize a photo to make it look as if it were drawn by a famous artist](#) and even [create photo-realistic faces](#) – not to mention create so-called [deepfakes](#) – the ability to infer properties such as [human emotions](#) from a real image has always been of interest for researchers.

Emotions are difficult because they tend to depend on context. For instance, when someone is concentrating on something it [might appear that they're simply thinking](#). [Facial recognition has come a long way](#) using machine learning, but identifying a person's [emotional](#) state based purely on looking at a person's face is missing key information.

Emotions are expressed not only through a person's expression but also

where they are and what they're doing. These contextual cues are difficult to feed into even modern machine learning algorithms. To address this, there are active efforts to [augment artificial intelligence techniques to consider context](#), not just for emotion recognition but all kinds of applications.

Reading employee emotions

The [report released by AI Now](#) sheds light on some ways in which AI is being applied to the workforce in order to evaluate worker productivity and even as early as at the interview stage. Analyzing footage from interviews, especially for remote job-seekers, [is already underway](#). If managers can get a sense of their subordinates' emotions from interview to evaluation, decision-making regarding other employment matters such as raises, promotions or assignments might end up being influenced by that information. But there are many other ways that this technology could be used.

Why the worry

These types of systems almost always have fairness, accountability, transparency and ethical ("FATE") flaws [baked into their pattern-matching](#). For example, one study found that facial recognition algorithms rated faces of black people as angrier than white faces, even when they were smiling.

[Many research groups are tackling this problem](#) but it seems clear at this point that the problem can't be solved exclusively at the technological level. Issues regarding FATE in AI will require a continued and concerted effort on the part of those using the technology to be aware of these issues and to address them. [As the AI Now report highlights](#): "Despite the increase in AI ethics content ... [ethical principles](#) and statements rarely focus on how AI ethics can be implemented and

whether they're effective." It notes that such AI ethics statements largely ignore questions of how, where, and who will put such guidelines into operation. In reality, it's likely that everyone must be aware of the types of biases and weaknesses these systems present, similar to how we must be aware of our own biases and those of others.

The problem with blanket technology bans

Greater accuracy and ease in persistent monitoring bring along other concerns beyond ethics. There are also a host of general technology-related privacy concerns, spanning from the [proliferation of cameras that serve as police feeds](#) to [potentially making sensitive data anonymous](#).

With these ethical and privacy concerns, a natural reaction might be to call for a ban on these techniques. Certainly, applying AI to job interview results or criminal sentencing procedures seems dangerous if the systems are learning biases or are otherwise unreliable. There are useful applications however, for instance in [helping spot warning signs to prevent youth suicide](#) and [detecting drunk drivers](#). That's one reason why even concerned researchers, regulators and citizens have generally stopped short of calling for blanket bans on AI-related technologies.

Combining AI and human judgment

Ultimately, technology designers and society as a whole need to look carefully at how information from AI systems is injected into decision-making processes. These systems can give incorrect results just like any other form of intelligence. [They are also notoriously bad](#) at rating their own confidence, not unlike humans, even in simpler tasks like [the ability to recognize objects](#). There also remain significant technical challenges in reading emotions, notably [considering context to infer emotions](#).

If people rely on a system that isn't accurate in making decisions, the users of that system are worse off. It's also well-known that [humans tend to trust these systems more than other authority figures](#). In light of this, we as a society need to carefully consider these systems' fairness, accountability, transparency and ethics both during design and application, always keeping a human as the final decision-maker.

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

Citation: AI can now read emotions—but should it? (2020, January 8) retrieved 16 April 2024 from <https://techxplore.com/news/2020-01-ai-emotionsbut.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.