An A.I. Pioneer Wants an FDA for Facial Recognition

Dave Gershgorn

Erik Learned-Miller helped create one of the most important face datasets in the world. Now he wants to regulate his creation.

Image: University of Massachusetts, Amherst

Erik Learned-Miller is one reason we talk about facial recognition at all.

In 2007, years before the current A.I. boom made “deep learning” and “neural networks” common phrases in Silicon Valley, Learned-Miller and three colleagues at the University of Massachusetts Amherst released a dataset of faces titled Labelled Faces in the Wild.

To you or me, Labelled Faces in the Wild just looks like folders of unremarkable images. You can download them and look for yourself. There’s a picture of Alec Baldwin pointing at someone off camera. There’s Halle Berry smiling at the Oscars. There’s boxer Joe Gatti, gloves raised mid-fight. But to an artificial intelligence algorithm, these folders contain the essence of what it means to look like a human.

This is why Labelled Faces in the Wild, often abbreviated as LFW, is so important. It’s cited in some of the most impactful facial recognition research in the past decade. When Google and Facebook were competing in 2014 and 2015 on facial recognition accuracy, the common test was performance categorizing LFW images, which exist in an unchanging database. LFW has been cited nearly 3,500 times, by researchers at Microsoft and Stanford, computer scientists in China and Hong Kong, and by Geoff Hinton, the guy responsible for neural networks in the first place.

It’s a big deal.

But now that artificial intelligence is also big business, Learned-Miller is thinking about ways to rein in the technology. His prevailing idea right now: Regulate A.I. the way the FDA regulates the medical

device industry. He doesn't have an official name for it yet, but one idea he's mulling: the FDA 2: the Facial recognition and Detection Agency.

> Learned-Miller is thinking about ways to rein in the technology. His prevailing idea right now: Regulate A.I. the way the FDA regulates the medical device industry.

One problem is that facial recognition doesn't work as well as it should. Research has shown that facial recognition used in the real world has problems with bias: It performs significantly better on people with lighter skin than people with darker skin and is better classifying males than females. This is attributed to biases in the data that algorithms analyze to learn the differences between faces. LFW, for instance, features mostly white males, so it's not surprising that algorithms trained on the dataset have trouble with faces that fall outside those parameters.

The technology landscape in 2019 is a lot different than it was in 2007, when LFW was first released. Artificial intelligence as we know it today was mostly research, and only a few schools, like NYU and the University of Toronto, were interested in so-called neural networks. Today, artificial intelligence can and is being used in the real world. Internet-connected cameras can send images to data centers to use facial recognition in real time, and smartphones — like newer iPhones with FaceID — commonly use the technology as a security feature. Government agencies have taken a keen interest in facial recognition as well. The FBI has used facial recognition based on driver's license and passport photos for law enforcement with little oversight for years, which Congress recently questioned during a House Oversight Committee hearing. What was research 10 years ago can now be carried in the palm of your hand.

For Learned-Miller, that research started in the 1980s and '90s from a simple idea: Humans are really, really good at recognizing faces. Replicating that ability would be an important step for bettering artificial intelligence as a whole, since the human skill of recognizing millions of different face patterns could be used for other computer vision tasks.

"From a science point of view, people's ability in face recognition has until very recently been considered just mind-blowingly good," Learned-Miller tells *OneZero.* "It's only our massive exposure to human faces, and the importance of recognizing properly, that causes us to have evolved and learn to differentiate amongst them. It's an incredibly interesting capability."

Learned-Miller tells a story about how, when walking down the streets of San Francisco, he recognized someone he hadn't seen in 25 years. Even after their face had changed due to age, and even though Learned-Miller had seen millions of faces seen in the years between, his brain was still able to recognize the person and remember who it was.

He says computer facial recognition just didn't work anywhere close to human level in 2007. Before bias was even an issue, researchers just wanted to make facial recognition work at all.

Inaccurate facial recognition can lead to problems of varying severity. On the milder end, perhaps a certain smartphone won't recognize darker-skinned faces as well for features like portrait lighting. On the other end, a product like Google Photos might associate black people with the keyword "gorilla." Or poor facial recognition could lead to law enforcement identifying the wrong person as a suspect in a crime.

This puts Learned-Miller in an uncomfortable position: The technology is already out in the world, but not the way it should be. He says he's been offered grants to produce another big facial recognition database but turned them down in order to continue research on how to build facial recognition datasets that are more equitable and compliant with new data laws like Europe's GDPR.

"You might say the easy way out is to stop doing research, but that's actually not an easy way out anymore, because we have unfair [datasets] out there," Learned-Miller says. "We're in this crazy kind of bind where it looks like there's no good direction to go, but we're not in a good place either."

Learned-Miller's FDA 2 solution isn't to abandon the research, but to regulate how it's used. His idea is to emulate the FDA's clearance process for medical devices. Since the limitations of medical devices can literally mean the difference between life and death, the FDA mandates exhaustive tests on the situations in which everything from catheters to surgical tools will and won't work. It's called 510(k) clearance.

"I used to be in the medical device industry, so I've authored these things myself, and they're huge documents that give you results of studies and all these other things," Learned-Miller says. "There's so many aspects of this which are a good fit with face recognition. One of the big ones is intended use and the data that you have support that intended use.

"And the opposite is, of course, the counter-indications: 'Look, we've never tested this software on nighttime images; therefore, you should not use the software on nighttime. We never tested this with children under 15, and therefore you shouldn't use it for children under 15,' and so forth," he says.

Labelled Faces in the Wild isn't going anywhere, and it's not changing. It can't — one of the most important aspects of LFW is that it stays constant. It's a benchmark, one where people can test their algorithms on the same level playing field.

But it could use disclosures. Learned-Miller says the dataset wasn't intended to be employed as a test to clear an algorithm for use in the real world.

> Just as a drug might come with a warning… so could a facial recognition system be labeled with do's and don'ts that clearly outline what users can expect.

"I may do this soon, put up a little disclaimer on the LFW site that says, 'Look, just because you did well on this database doesn't mean that your software is ready for deployment," he says. "This is not a golden check mark that says your face recognition is ready for anything. I think most people understand that, but some don't. It doesn't look at kids. It doesn't look at really old people. There are not that many women in it."

This brings us back to Learned-Miller's FDA 2 idea. The FDA exists to ensure a supplement on the drugstore shelf won't kill you, but there's no such oversight for facial recognition at the moment, even though it, too, could have fatal consequences if used incorrectly. Just as a drug might come with a warning that those taking it shouldn't drink or work heavy machinery, so could a facial recognition system be labeled with clear do's and don'ts that outline what users can expect.

"These are all big standard mechanisms that dramatically improve the efficacy and safety of a drug," Learned-Miller says. "Nobody would think about going back to the days when we didn't have [the FDA]. So I think it's doable, but it's going to take time."