LOS ANGELES & SAN FRANCISCO
Daily Journal
www.dailyjournal.com
WEDNESDAY, JUNE 26, 2019

TOP ARTIFICIAL INTELLIGENCE LAWYERS

Before We Regulate

By H. Mark Lyon

Each day there are new stories in the press about how machine learning and other "artificial intelligence" technologies are changing the world, sometimes for the better, sometimes not. Increasingly, many of these stories discuss whether to regulate AI to ensure that it is put to beneficial use. While there are already a number of laws and regulations that apply to AI products and services, there are still relatively few — at least in the U.S. — directed specifically to the question of when and how we will allow AI technologies to be used in our lives. For the most part, the U.S. (both at the federal and state levels) has taken a "wait and see" approach, allowing AI technologies and their use cases to develop at their own pace.

But with increasingly alarmist and strident headlines appearing almost daily about some of AI's perceived failures, change may well be on the horizon. Last month, San Francisco passed one of the strictest U.S. laws to date governing machine learning and AI technologies. The Stop Secret Surveillance ordinance amends San Francisco's Administrative Code to require that city departments seeking to acquire or use surveillance technologies must first submit certain reports and seek certain purchase and use approvals from the San Francisco Board of Supervisors before any purchase or use of such technologies. The purpose behind the ordinance was to eliminate the possibility that law enforcement or other public officials would secretly utilize surveillance technology, without having first run it by the board and its available public hearing mechanisms. However, with regard to one machine learning-driven technology, facial recognition, the ordinance goes one step further and completely bans any use of any form of the technology by all city and county departments, including the San Francisco Police Department. Thus, despite the potential for targeted, and important, uses of the technology, such as assisting with securing public buildings and spaces (either generally or during the limited period of a known threat), San Francisco chose to simply ban all uses of facial recognition technology in any form.

H. Mark Lyon *is partner at Gibson, Dunn & Crutcher LLP and is chair of the firm's Artificial Intelligence and Automated Systems Practice Group.*

An absolute ban on all uses of all forms of a particular technology is rarely, perhaps even never, appropriate, at least where less extreme approaches are available. For technology, use and form matter. Technology itself is neither inherently good nor bad; rather, it is the uses we make of the technology, or the particular forms of the technology being used, that may be problematic. Before we attempt to regulate technology, we need to understand how the uses and forms of the technology may contribute to the perceived problems. Any regulation should then be carefully framed to address the problem, based on a reasoned, factual inquiry of all available evidence, and tailored to leave room for later resolution through potential improvements in the technology.

Importantly, there isn't just one use case for facial recognition technology, there are numerous possible use cases. Ask a parent whose child is missing, whether they would welcome the use of facial recognition technology to help find a lost or abducted child in a crowd. Only those who stand on the most rigid of principles will continue to argue that limited applications of facial recognition in such urgent circumstances are still more harmful than beneficial. And yet, it is of course true that some potential uses, such as Western China's use of facial recognition technology to profile a Muslim minority group for potential discrimination, go against our strong values of personal liberty and would likely not be permissible, with or without any ban. In addition, as with all technologies, there are different facial recognition product offerings available from different purveyors that each perform with differing capabilities and limitations. There is no single form of technology that is "facial recognition." As a result, even were

Despite the potential for targeted, and important, uses of the technology, such as assisting with securing public buildings and spaces (either generally or during the limited period of a known threat), San Francisco chose to simply ban all uses of facial recognition technology in any form.

There is no single form of technology that is 'facial recognition.' As a result, even were we to agree with supporters of an outright ban on facial recognition that certain products or forms of the technology may have flaws or limitations that make them undesirable, other forms of the technology — perhaps future and more refined product releases — may instead perform acceptably for their intended uses.

we to agree with supporters of an outright ban on facial recognition that certain products or forms of the technology may have flaws or limitations that make them undesirable, other forms of the technology — perhaps future and more refined product releases — may instead perform acceptably for their intended uses. Again, use and form matter.

But the San Francisco Board of Supervisors appears to have been overly influenced both by the repressive actions of China and from studies reported by the ACLU and MIT, which concluded that Amazon's "Rekognition" facial recognition software disproportionately misidentified minorities and women. Indeed, the ordinance's supporter, Supervisor Aaron Peskin, was quoted in the S.F. Examiner as saying that it is a "fact" that facial recognition technology "has the biases of the people who developed it." Thus, the board adopted an ordinance that asserts, as a factual matter without meaningful factual support, that facial recognition technology "will exacerbate racial injustice." (Emphasis added.)

However, such blanket statements are simply not true. Not all forms of facial recognition software will necessarily have all the biases of everyone who developed the software. Amazon's Rekognition software is not the sole example of all possible facial recognition software; rather it is just one form. If, as the ACLU and MIT assert, the Rekognition software is flawed and biased, then the answer is to ensure that the software is improved to eliminate those flaws and biases, and to consider its acceptable uses, before such flaws and biases can prove harmful. Indeed, there are currently existing techniques to identify and eliminate or correct unwanted biases in machine-learning software, and these techniques rapidly continue to improve.

And importantly, it is not entirely clear that the ACLU and MIT were correct in their assessment that Amazon's software is as flawed as claimed. Responding to their studies, Amazon has noted that confidence levels of 80% (the levels used by both ACLU and MIT studies) are not appropriate in the law enforcement context, and that higher confidence levels (perhaps as high as 99%) would have produced appropriate results. It is thus possible that the San Francisco Board of Supervisor's concerns with reported bias might have been alleviated simply by increasing the required confidence level in the identification made by the facial recognition software, or by requiring some other corroboration of outputs to achieve a sufficiently accurate response.

In either case, the board's outright ban on facial recognition technology was shortsighted and an overreaction. A better response, which would still have achieved the desired goals, would have taken a less extreme approach. For example, as it did for other forms of surveillance technology, the board could have elected to require board review and approval of all proposed uses of facial recognition to allow the board to consider the impact of bias specific to each proposed use case and impose suitable limits on the circumstances and scope of any approved use. While it is possible that existing facial recognition software would still be found lacking, perhaps later, improved versions, or even some other company's facial recognition solution, would prove worthy for use under critical circumstances such as searching for missing or abducted children, or heightened security for public locations or schools in times of known threats.

We should feel free to regulate AI-based technologies when appropriate to avoid harmful impacts. However, any such regulations need to be thoughtful and carefully considered, otherwise we risk the baby as we toss out the bath water.

The views, thoughts and opinions expressed herein are those of the author only and do not reflect the views of Gibson, Dunn & Crutcher LLP or any of its clients.