

ethikos Volume 34, Number 4. April 01, 2020

What artificial intelligence bias can teach us about holistic product design

By Dylan Doyle-Burke

Dylan Doyle-Burke is CEO of Radical AI LLC in Denver, CO.

twitter.com/dylandoyleburke

Since its most recent boom in 2012, artificial intelligence (AI) has captured the imaginations of corporate industry and society at large. Technical developments in the field have evolved quickly, with companies racing to corner the market on exciting new technology and capability. Unfortunately, as a side effect of this speed of progress, the pace of technical development soon largely outpaced the ethical conversations needed to provide intentional moral structure to AI product design.

Particularly on the issues of privacy, fairness, bias, accountability, and transparency, the rapid development of AI technology has caused the field of AI technology development to fall into several high-profile ethical crises. The most notable of these crises was a ProPublica feature on machine learning software used in police systems across the country to assign a recidivism value to those who were incarcerated that would then be used by a judge in sentencing. Through presenting rigorous evidence, ProPublica argued that the prediction of future criminality that such software provided was largely incorrect and was definitively biased against black offenders.^[1] The ProPublica feature spread through the media zeitgeist, and the field of AI was forced to begin to reflect on what to do with its bias problem.

AI bias

The problem of machine learning bias was just beginning. Soon after the ProPublica piece was published, the University of Toronto's Inioluwa Deborah Raji and MIT's Joy Buolamwini released a study testing the results of facial recognition technology from two major US tech giants, Microsoft and IBM, and a Chinese AI company, Face++. Raji and Buolamwini found that every instance of facial recognition technology they tested performed better for lighter-skinned faces than for darker-skinned faces. The reason for this, they concluded, was because of the content of the training data used for the facial recognition machines.^[2] Their study highlighting blatant AI bias was again picked up by the media in a series of articles skeptical about the ethics inherent in machine learning.

Examples of bias and other fairness concerns regarding AI technology development continued to surface. Studies on machine learning algorithms used in teacher evaluation, automated hiring tools, automated advertisement selection, and more revealed that the field of AI technology development necessitated some sort of strategy to increase fairness and decrease its bias. Questions of access, data selection, data implementation, product design, and more were raised within the industry and the academy alike. Many different strategies were attempted at different levels of the AI product development cycle, and many were successful.

However, what quickly became clear is that to truly decrease bias and increase the chance of releasing an ethical AI product, companies would need to implement strategies throughout the entire product development cycle, not

just at one or two points throughout it. The goal has become to alter the entire product development cycle in order to create a more ethical product. To do this, effective AI technology development corporations have begun to holistically alter their product development systems from problem definition to quality assurance and customer interaction.

This document is only available to subscribers. Please [log in](#) or [purchase access](#).

[Purchase Login](#)