

2020 Research Paper Competition

Biased Programmers? Or Biased Data? A Field Experiment in Operationalizing

AI Ethics

Fabrizio Dell'Acqua, Bo Cowgill, Samuel Deng, Daniel Hsu, Nakul Verma,

Augustin Chaintreau

Columbia Business School and Columbia University Department of Computer
Science

Why do biased beliefs arise, and what interventions can prevent them? We quantify the effects of these using a field experiment using machine learning to predict human capital. We ask 400 AI engineers to predict standardized test scores for a representative sample of OECD residents. Engineers were randomly assigned to four treatment conditions affecting their training data, awareness of bias issues and understanding of bias-reduction technical methods. We then assess the resulting predictive algorithms using the realized test performance and through randomized audit-like manipulations of algorithmic inputs.