

Letter from the Editor-in-Chief

Over the past few years, we have come to realize that a great danger of technology is the unfair bias that AI systems create, reinforce, and propagate. While the source of such bias is usually human, machines' homogeneity and efficiency could easily magnify its damage, a concern further exacerbated by the fact that AI systems are being trusted to make more and more critical decisions in almost every aspect of human experience.

Data is often the channel through which bias transmits from human experience to machines. For example, a training dataset with inherent societal bias, or simply a limited and unrepresentative training dataset, may lead to machine learning models that are unfair and biased towards a specific class. However, detecting and correcting unfair bias in AI systems are much more than just a data management task. In particular, since bias starts with humans and will eventually affect humans, addressing unfair bias requires insights from psychological, societal, and cultural perspectives.

James Foulds and Shimei Pan put together the current issue – Interdisciplinary Perspectives on Fairness and Artificial Intelligence Systems – that consists of six papers from leading researchers in multidisciplinary areas. It is an excellent collection of ideas and best practices on this important topic.

Haixun Wang
Instacart