

Letter from the Editor-in-Chief

A year ago, we devoted a special issue of the Data Engineering Bulletin to the state-of-the-art work of understanding the different types of bias that is created, reinforced, and propagated by AI systems.

We would like to dive into a broader topic, Ethics of Artificial Intelligence, which is concerned with how humans interact with (design, use, and treat) AI systems in general. Shimei Pan and James Foulds put together the current issue—Responsible AI and Human-AI Interaction—to provide an overview of the recent efforts in this domain. The issue consists of five papers from leading researchers in the AI and the Human-Computer Interaction communities, and covers topics such as how people perceive AI, how to improve AI’s transparency and interpretability, and how to foster effective collaborations between humans and machines.

Unlike our typical issues of the Data Engineering Bulletin that mostly focus on data management or data-driven challenges and solutions, this issue brings up and highlights human-centered approaches, which we believe are of increasing importance in the new era of computing. The human-centered initiative provides a new angle for us to consider questions ranging from whether AI is a threat to human existence or a promise for humans to achieve unprecedented levels of creativity and productivity to what is the role of Ethics of AI in the broad field of artificial intelligence and computing.

Haixun Wang
Instacart