

Podcast Name: *ACM ByteCast*

Episode: Kush Varshney - Episode 43

Welcome to the *ACM ByteCast* podcast, a series from the Association for Computing Machinery! The podcast features conversations with researchers, practitioners, and innovators at the intersection of computing research and practice about their experiences, lessons learned, and visions for the future of computing. In this episode, host Bruke Kifle interviews guest Dr. Kush Varshney about responsible AI and how it can be achieved.

Dr. Kush Varshney is a distinguished research scientist and manager at IBM Research in New York. He leads the machine learning group in the Foundations of Trustworthy AI Department, where he applies data science and predictive analytics to the fields of healthcare, public affairs, algorithmic fairness and international development. He is also the founding co-director of the IBM Science for Social Good Initiative. He has contributed to the development of several open-source toolkits such as AI Fairness 360 and AI Explainability 360. In 2022, he independently published the book *Trustworthy Machine Learning*.

To begin, Dr. Varshney shares a few key moments which have helped to shape the course of his career thus far. He became progressively more interested in machine learning during his time in graduate school at MIT. He knew he wanted to do research, and IBM made a lot of sense for him. Revolutionary at the time, the group he joined was using machine learning to make predictions in various contexts. Several years later he began working with DataKind, an organization connecting practice and data scientists with nonprofits and social change organizations to do applied work. Today, the organization has done many projects addressing things like education, inequality, world hunger and poverty.

Dr. Varshney defines responsible AI as an overarching umbrella which encompasses AI ethics, trustworthy AI and AI governance. It is crucial for AI systems to be trustworthy, reliable, and fair. Dr. Varshney also strongly believes in empowering everyone to use AI technologies for meeting their own goals and purposes. Then, he unpacks the idea of finding balance between the technical and social aspects of AI. The tools for operationalizing RAI are a starting point, not the entire solution. The motivation behind open sourcing was to bridge the gap between academic/industrial labs and practice. Dr. Varshney unpacks the inherent risks of these emerging AI technologies. When the output becomes generative content, there are even more new risks which can emerge.

Next, the conversation shifts to discuss the call for pausing the development of new AI systems. Now that we have the first 50 years of AI behind us, the main focus should be on safety. Regulations don't just apply to laws, but also to shifting social norms through the market. Dr. Varshney believes it will be a natural progression of putting safety first in the field as we look towards the future. While AI does pose certain major risks to humanity, there is also a concern around its small and tangible short-term risks. All stakeholders have a role to play in the path towards responsible AI coordination. Everyone should be bringing forth their own principles and policies, and systems should be designed in a way that the input guides and controls the

behavior of the system. Government regulation, self-regulation and watchdog civil society organizations all have their own place in this.

Before wrapping up, Dr. Varshney shares the inspiration behind his own self-published book *Trustworthy Machine Learning*, which was very much a passion project for him. He began writing the book before Covid during a time when he thought people really needed to hear the unique experience he had to share. He decided to self-publish in order to keep this critical information as accessible as possible for everyone. The biggest takeaway he hopes readers get from the book is that taking shortcuts is not very responsible. Then, he touches on his work with IBM's Science for Social Good Initiative. Finally, Dr. Varshney shares what he is most excited about in regard to the future of AI.

Key takeaways:

2:17 - Dr. Varshney shares the milestone moments over the course of his professional career.

10:40 - What is responsible AI?

17:04 - Balancing the technical and social aspects of AI.

23:00 - Operationalizing RAI principles.

28:24 - The inherent risks of these emerging technologies.

31:04 - The call for a pause of the development of new AI systems.

36:33 - The role of stakeholders in the path towards AI coordination.

38:06 - The inspiration behind Dr. Varshney's book *Trustworthy Machine Learning*.

43:22 - Dr. Varshney's work with IBM's Science for Social Good Initiative.

51:02 - Exciting evolutions in AI.

Links

Learn more about [Kush Varshney](#).

Learn more about [Bruke Kifle](#).

Learn more about the ACM ByteCast podcast at <https://learning.acm.org/bytecast>

Tags:

Kush Varshney, AI, artificial intelligence, generative technology, chatGPT, machine learning, data science, predictive analytics, open source, MIT, IBM, social impact, technology, human capital, data scientists, AI ethics, AI governance, safe technology, digitization, generative AI