In research, AI is increasingly used in predictive-modeling scenarios where the algorithm learns from data to predict a future outcome. Much of this work relies on machine learning, a branch of AI that, as Hao puts it, trains "algorithm[s] to understand often massive amounts of data, find the meaningful patterns in that data, and then apply those patterns to future data."

One challenge in evaluating AI studies is their "black box" nature such that researchers only see what goes into the algorithm and what comes out. As Hutson and others have reported, the mystery of what the algorithm does in between has contributed to a reproducibility crisis in AI and highlights the peril of overgeneralizing the results of any single AI study. For journalists, it's essential to stringently assess the results and how they compare to what's been done before. To do so, it can be especially useful to reach out to AI researchers not involved in the new work to get their take on whether the results are really as novel as they're claimed to be.

Finally, check whether a new AI model offers a significant and realistic advance before covering it. It's useful to ask whether the AI model is being compared to a valid benchmark, such as the current state-of-the-art in the field, Hutson says. Hao also talks to experts in the field who aren't using AI themselves. For a study claiming that its deep learning algorithm could help improve education outcomes, for example, Hao might ask an education specialist whether they think the AI model could actually be implemented in a real-world setting.

"Oftentimes, as science writers, we don't question the usefulness of technology. Do we need this? That's always something that's worth asking," McCullom says.

Commented [1]: For examples in a variety of research fields, see:
healthcare: https://www.nature.com/articles/s41746-021-00549-7

engineering: https://www.sciencedirect.com/science/article/abs/pii/S0029801821005680

energy use: https://www.sciencedirect.com/science/article/abs/pii/S1364032116307420

Commented [2]: see first paragraph in https://www.ibm.com/topics/machine-learning

Commented [3]: Confirm with Hao and see Hao transcript and audio at 3:14

Commented [4]: Several sources mentioned this black box issue — confirm with Hao, Hutson, and McCullom

Commented [5]: Confirm with Hutson

Commented [6]: Confirm with Hutson and Hao

Commented [7]: This particular suggestion came from my own expertise in reporting on AI rather than interviews with sources.

Commented [8]: Confirm with Hutson

Commented [9]: Confirm with Hao

Commented [10]: Confirm with McCullom and see McCullom transcript and audio at 45:31