

Deepfakes aren't very good—nor are the tools to detect them | Ars Technica

By Will Knight

June 15, 2020

We're lucky that [deepfake videos](#) aren't a big problem yet. The best deepfake detector to emerge from a major Facebook-led effort to combat the altered videos would only catch about two-thirds of them.

In September, as speculation about the danger of deepfakes grew, [Facebook](#) challenged [artificial intelligence](#) wizards to develop techniques for detecting deepfake videos. In January, the company also [banned deepfakes](#) used to spread misinformation.

Facebook's [Deepfake Detection Challenge](#), in collaboration with Microsoft, Amazon Web Services, and the [Partnership on AI](#), was [run through Kaggle](#), a platform for coding contests that is owned by Google. It provided a vast collection of face-swap videos: 100,000 deepfake clips, created by Facebook using paid actors, on which entrants tested their detection algorithms. The project attracted more than 2,000 participants from industry and academia, and it generated more than 35,000 deepfake detection models.

[...]

Source: [Deepfakes aren't very good—nor are the tools to detect them | Ars Technica](#)