

Research reflects how AI sees through the looking glass

July 2 2020



AI learns to pick up on unexpected clues to differentiate original images from their reflections, the researchers found. Credit: Cornell University

Things are different on the other side of the mirror.

Text is backward. Clocks run counterclockwise. Cars drive on the wrong side of the road. Right hands become left hands.

Intrigued by how reflection changes images in subtle and not-so-subtle ways, a team of Cornell University researchers used artificial intelligence to investigate what sets originals apart from their reflections. Their algorithms learned to pick up on unexpected clues such as hair parts, [gaze direction](#) and, surprisingly, beards—findings with implications for training [machine learning](#) models and detecting faked images.

"The universe is not symmetrical. If you flip an image, there are differences," said Noah Snavely, associate professor of computer science at Cornell Tech and senior author of the study, "Visual Chirality," [presented](#) at the 2020 Conference on Computer Vision and Pattern Recognition, held virtually June 14-19. "I'm intrigued by the discoveries you can make with new ways of gleaning information."

Zhiqui Lin is the paper's first author; co-authors are Abe Davis, assistant professor of computer science, and Cornell Tech postdoctoral researcher Jin Sun.

Differentiating between original images and reflections is a surprisingly easy task for AI, Snavely said—a basic deep learning [algorithm](#) can quickly learn how to classify if an image has been flipped with 60% to 90% accuracy, depending on the kinds of images used to train the algorithm. Many of the clues it picks up on are difficult for humans to notice.

For this study, the team developed technology to create a heat map that indicates the parts of the image that are of interest to the algorithm, to gain insight into how it makes these decisions.

They discovered, not surprisingly, that the most commonly used clue was text, which looks different backward in every written language. To learn more, they removed images with text from their data set, and found that

the next set of characteristics the model focused on included wrist watches, shirt collars (buttons tend to be on the left side), faces and phones—which most people tend to carry in their right hands—as well as other factors revealing right-handedness.

The researchers were intrigued by the algorithm's tendency to focus on faces, which don't seem obviously asymmetrical. "In some ways, it left more questions than answers," Snavely said.

They then conducted another study focusing on faces and found that the heat map lit up on areas including hair part, eye gaze—most people, for reasons the researchers don't know, gaze to the left in portrait photos—and beards.

Snavely said he and his team members have no idea what information the algorithm is finding in beards, but they hypothesized that the way people comb or shave their faces could reveal handedness.

"It's a form of visual discovery," Snavely said. "If you can run machine learning at scale on millions and millions of images, maybe you can start to discover new facts about the world."

Each of these clues individually may be unreliable, but the algorithm can build greater confidence by combining multiple clues, the findings showed. The researchers also found that the algorithm uses low-level signals, stemming from the way cameras process images, to make its decisions.

Though more study is needed, the findings could impact the way machine learning models are trained. These models need vast numbers of images in order to learn how to classify and identify pictures, so computer scientists often use reflections of existing images to effectively double their datasets.

Examining how these reflected images differ from the originals could reveal information about possible biases in machine learning that might lead to inaccurate results, Snively said.

"This leads to an open question for the computer vision community, which is, when is it OK to do this flipping to augment your dataset, and when is it not OK?" he said. "I'm hoping this will get people to think more about these questions and start to develop tools to understand how it's biasing the algorithm."

Understanding how reflection changes an image could also help use AI to identify images that have been faked or doctored—an issue of growing concern on the internet.

"This is perhaps a new tool or insight that can be used in the universe of image forensics, if you want to tell if something is real or not," Snively said.

More information: linzhiqiu.github.io/papers/chirality/main.pdf

Provided by Cornell University

Citation: Research reflects how AI sees through the looking glass (2020, July 2) retrieved 23 April 2024 from <https://techxplore.com/news/2020-07-ai-glass.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--