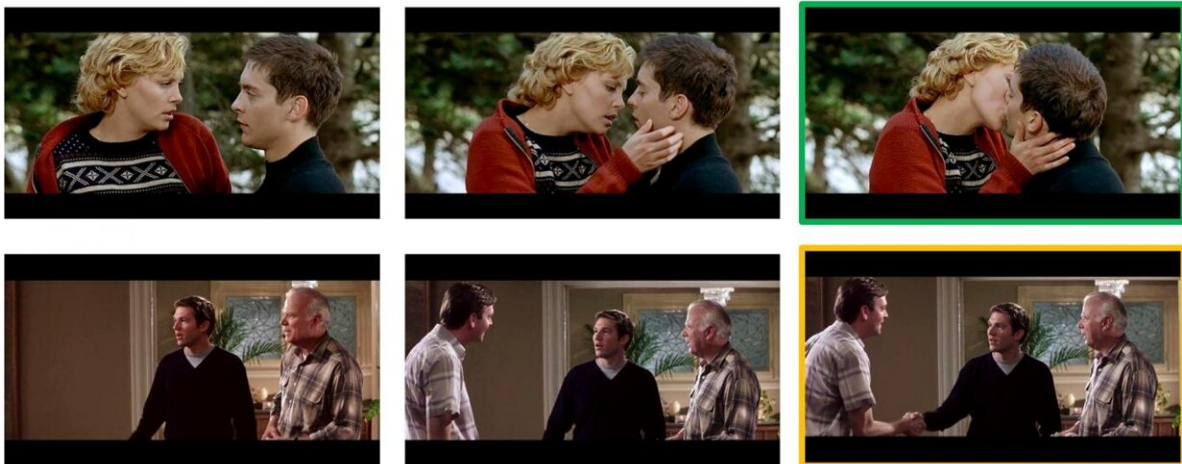


AI learns to predict human behavior from videos

June 28 2021



The AI model recognizes when the future is uncertain and is capable of "hedging the bet," the way a person would, accordingly. For instance, when the model finds it impossible to predict whether two people are going to hug or handshake, it predicts they are going to greet each other instead. Credit: Dídac Surís/Columbia Engineering

Predicting what someone is about to do next based on their body language comes naturally to humans but not so for computers. When we meet another person, they might greet us with a hello, handshake, or even a fist bump. We may not know which gesture will be used, but we can read the situation and respond appropriately.

In a new study, Columbia Engineering researchers unveil a [computer vision](#) technique for giving [machines](#) a more intuitive sense for what will happen next by leveraging higher-level associations between people, animals, and objects.

"Our algorithm is a step toward machines being able to make better predictions about [human behavior](#), and thus better coordinate their actions with ours," said Carl Vondrick, assistant professor of computer science at Columbia, who directed the study, which was presented at the International Conference on Computer Vision and Pattern Recognition on June 24, 2021. "Our results open a number of possibilities for human-robot collaboration, autonomous vehicles, and assistive technology."

It's the most accurate method to date for predicting video action events up to several minutes in the future, the researchers say. After analyzing thousands of hours of movies, sports games, and shows like "The Office," the system learns to predict hundreds of activities, from handshaking to fist bumping. When it can't predict the specific action, it finds the higher-level concept that links them, in this case, the word "greeting."

Past attempts in predictive machine learning, including those by the team, have focused on predicting just one action at a time. The algorithms decide whether to classify the action as a hug, high five, handshake, or even a non-action like "ignore." But when the uncertainty is high, most machine learning models are unable to find commonalities between the possible options.

Columbia Engineering Ph.D. students Didac Suris and Ruoshi Liu decided to look at the longer-range prediction problem from a different angle. "Not everything in the future is predictable," said Suris, co-lead author of the paper. "When a person cannot foresee exactly what will happen, they play it safe and predict at a higher level of abstraction. Our

algorithm is the first to learn this capability to reason abstractly about future events."

Suris and Liu had to revisit questions in mathematics that date back to the ancient Greeks. In high school, students learn the familiar and intuitive rules of geometry—that straight lines go straight, that parallel lines never cross. Most machine learning systems also obey these rules. But other geometries, however, have bizarre, counter-intuitive properties; straight lines bend and triangles bulge. Suris and Liu used these unusual geometries to build AI models that organize high-level concepts and predict human behavior in the future.

"Prediction is the basis of human intelligence," said Aude Oliva, senior research scientist at the Massachusetts Institute of Technology and co-director of the MIT-IBM Watson AI Lab, an expert in AI and human cognition who was not involved in the study. "Machines make mistakes that humans never would because they lack our ability to reason abstractly. This work is a pivotal step towards bridging this technological gap."

The mathematical framework developed by the researchers enables machines to organize events by how predictable they are in the future. For example, we know that swimming and running are both forms of exercising. The new technique learns how to categorize these activities on its own. The system is aware of uncertainty, providing more specific actions when there is certainty, and more generic predictions when there is not.

The technique could move computers closer to being able to size up a situation and make a nuanced decision, instead of a pre-programmed action, the researchers say. It's a critical step in building trust between humans and computers, said Liu, co-lead author of the paper. "Trust comes from the feeling that the robot really understands people," he

explained. "If machines can understand and anticipate our behaviors, computers will be able to seamlessly assist people in daily activity."

While the new algorithm makes more accurate predictions on benchmark tasks than previous methods, the next steps are to verify that it works outside the lab, says Vondrick. If the system can work in diverse settings, there are many possibilities to deploy machines and robots that might improve our safety, health, and security, the researchers say. The group plans to continue improving the algorithm's performance with larger datasets and computers, and other forms of geometry.

"Human behavior is often surprising," Vondrick commented. "Our algorithms enable machines to better anticipate what they are going to do next."

The study is titled "Learning the predictability of the future."

More information: Dídac Surís et al, Learning the Predictability of the Future. arXiv:2101.01600 [cs.CV] arxiv.org/abs/2101.01600

PDF link: [openaccess.thecvf.com/content/... CVPR_2021_paper.pdf](https://openaccess.thecvf.com/content/CVPR2021/papers/Suris_Learning_the_Predictability_of_the_Future_Suris2021/CVPR_2021_paper.pdf)

Provided by Columbia University School of Engineering and Applied Science

Citation: AI learns to predict human behavior from videos (2021, June 28) retrieved 19 April 2024 from <https://techxplore.com/news/2021-06-ai-human-behavior-videos.html>

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.