

Artificial intelligence / Machine learning

How special relativity can help AI predict the future




MIT Technology Review

about what will happen next. It could help predict the effects of drugs.

by **Will Douglas Heaven**

August 28, 2020



ÉTIENNE-JULES MAREY

Nobody knows what will happen in the future, **but some guesses are a lot better than others. A kicked football will not reverse in midair and return to the kicker's foot. A half-eaten cheeseburger will not become whole again. A broken arm will not heal overnight.**

By drawing on a fundamental description of cause and effect found in Einstein's theory of special

relativity, researchers from Imperial College London have come up with a way to help AIs make better guesses too.

The world progresses step by step, every instant emerging from those that precede it. We can make good guesses about what happens next because we have strong intuitions about cause and effect, honed by observing how the world works from the moment we are born and processing those observations with brains hardwired by millions of years of evolution.



Sign up for **The Algorithm** - Artificial intelligence, demystified

Sign up

You've read **1 of 3**

[Sign in](#)

Subscribe

Computers, however, find causal reasoning hard. Machine-learning models excel at spotting correlations but are hard pressed to explain why one event should follow another. That's a problem because without a sense of cause and effect, predictions can be wildly off. Why shouldn't a football reverse in flight?

This is a particular concern with AI-powered diagnosis. Diseases are often correlated with multiple symptoms. For example, people with type 2 diabetes are often overweight and have shortness of breath. But the shortness of breath is not caused by the diabetes, and treating a patient with insulin will not help with that symptom.

The AI community is realizing how important causal reasoning could be for machine learning and is scrambling to find ways to bolt it on.

Researchers have tried various ways to help computers predict what might happen next. Existing approaches train a machine-learning model frame by frame to spot patterns in sequences of action. Show the AI a few frames of a train pulling out of a station and then ask it to generate the next few frames in the sequence, for example.

AIs can do a good job of predicting a few frames into the future, but the accuracy falls off sharply after five or 10 frames, says Athanasios Vrontzos at Imperial College London. Because the AI uses preceding frames to generate the next one in the sequence, small mistakes made early on—a few glitchy pixels, say—get compounded into larger errors as the sequence progresses.

Vrontzos and his colleagues wanted to try a different approach. Instead of getting an AI to learn to predict a specific sequence of future frames by watching millions of video clips, they allowed it to generate a whole range of frames that were roughly similar to the preceding ones and then pick the

that were most likely to come next. The AI can make guesses about the future without having to learn anything about the progression of time, says Vlontzos.

To do this, the team developed an algorithm inspired by light cones, a mathematical description of boundaries of cause and effect in spacetime, which was first proposed in Einstein's theory of special relativity and later refined by his former professor Hermann Minkowski. Light cones emerge in physics because the speed of light is constant. They show the expanding limits of a ray of light—anything else—as it emanates from an initial event, such as an explosion.

Take a sheet of paper and mark an event on it with a dot. Now draw a circle with that event at the center. The distance between the dot and the edge of the circle is the distance light has traveled in a period of time—say, one second. Because nothing, not even information, can travel faster than light, the edge of this circle is a hard boundary on the causal influence of the original event. In principle, anything inside the circle could have been affected by the event; anything outside could not.

After two seconds, light has traveled twice the distance and the circle's size has doubled: there are now many more possible futures for that original event. Picture these ever larger circles rising second by second out of the sheet of paper, and you have an upside-down cone with the original event at its tip. This is a light cone. A mirror image of the cone can also extend backwards, behind the sheet of paper; it will contain all possible pasts that could have led to the original event.

Vlontzos and his colleagues used this concept to constrain the future frames an AI could pick. They tested the idea on two data sets: Moving MNIST, which consists of short video clips of handwritten digits moving around on a screen, and the KTH human action series, which contains clips of people walking or waving their arms. In both cases, they trained the AI to generate frames that looked similar to those in the data set. But importantly the frames in the training data set were not shown in sequence, and the algorithm was not learning how to complete a series.

They then asked the AI to pick which of the new frames were more likely to follow another. To do this, the AI grouped generated frames by similarity and then used the light-cone algorithm to draw a boundary around those that could be causally related to the given frame. Despite not being trained to continue a sequence, the AI could still make good guesses about which frames came next. If you give the AI a frame in which a short-haired person wearing a shirt is walking, then the AI will reject frames that show a person with long hair or no shirt, says Vlontzos. The work is in the final stages of review at NeurIPS, a major machine-learning conference.

An advantage of the approach is that it should work with different types of machine learning, as long as the model can generate new frames that are similar to those in the training set. It could also be used to improve the accuracy of existing AIs trained on video sequences.

To test the approach, the team had the cones expand at a fixed rate. But in practice, this rate will vary

A ball on a football field will have more possible future positions than a ball traveling along rails, for example. This means you would need a cone that expanded at a faster rate for the football.

Working out these speeds involves getting deep into thermodynamics, which isn't practical. For now, the team plans to set the diameter of the cones by hand. But by watching video of a football game, the AI could learn how much and how fast objects moved around, which would enable it to set the diameter of the cone itself. An AI could also learn on the fly, observing how fast a real system changes and adjusting cone size to match it.

Predicting the future is important for many applications. Autonomous vehicles need to be able to predict whether a child is about to run into the road or whether a wobbling cyclist presents a hazard. Robots that need to interact with physical objects need to be able to predict how those objects will behave when moved around. Predictive systems in general will be more accurate if they can reason about cause and effect rather than just correlation.

But Vlontzos and his colleagues are particularly interested in medicine. An AI could be used to simulate how a patient might respond to a certain treatment—for example, spooling out how that treatment might run its course, step by step. “By creating all these possible outcomes, you can see how a drug will affect a disease,” says Vlontzos. The approach could also be used with medical images. Given an MRI scan of a brain, an AI could identify the likely ways a disease could progress.

“It’s very cool to see ideas from fundamental physics being borrowed to do this,” says Ciaran Lee, a researcher who works on causal inference at Spotify and University College London. “A grasp of causality is really important if you want to take actions or decisions in the real world,” he says. It goes to the heart of how things come to be the way they are: “If you ever want to ask the question ‘Why?’ then you need to understand cause and effect.”

Share       Link 

Tagged Artificial Intelligence, medicine

Author  [Will Douglas Heaven](#)

The coronavirus responders

A special series on coronavirus responses around the world.



01.
**How Mongolia has kept
the coronavirus at bay**
Aug 18

02.
**How Ebola helped Liberia
prepare for covid**
Aug 19

03.
No more maté sharing
Aug 19

Election 2020 10 hours

Explainer: What do political databases know about you?

If you live in the US, you're almost certainly being tracked by political organizations. They know a lot about you—but some data is just guesswork.

Elon Musk's Neuralink is neuroscience theater

Elon Musk's livestreamed brain implant event made promises that will be hard to keep.



Artificial intelligence 3 days

IBM has built a new drug-making lab entirely in the cloud

AI, robotics, and the cloud are allowing scientists to run chemistry experiments remotely.



Artificial intelligence 3 days

How special relativity can help AI predict the future

An approach inspired by Einstein's light cones lets AIs make better guesses about what will happen next. It could help predict the effects of drugs.

E-learning? There's
a database for that.
Real-time data?
That, too.

Learn more


In association with

aws databases

Artificial Intelligence

Welcome to In Machines We Trust

A podcast about the automation of everything



01.
**The scientists who swab
subways for coronavirus**
Aug 26

02.
**What happens when an
algorithm gets it wrong**
Aug 12

03.
Land of a billion faces
Aug 12

SPONSORED

Amid the covid-19 pandemic, shifting business priorities

Organizations are reshuffling projects and accelerating investments that were already underway, leaning heavily on technology to stay competitive.




In association with Inference Solutions

Artificial intelligence 3 days

Memers are making deepfakes, and things are getting weird

The rapidly increasing accessibility of the technology raises new concerns about its abuse.




Bottom right: it me.

Computing Aug 28

How a \$1 million plot to hack Tesla failed

A Tesla employee thwarted a ransomware attack that targeted the Gigafactory



Israeli phone hacking company faces court fight over sales to Hong Kong

...

E-learning? There's a database for that. Real-time data? That, too

...



New MIT Technology Review Podcast

Deep Tech

Our journalists discuss in depth the most important stories they are reporting, from covid-19 tracing to personal preparedness in a pandemic.

Hear the **subscriber-only** podcast.

[Listen now](#)

Load more