

Helping machines perceive some laws of physics

December 2 2019, by Rob Matheson



An MIT-invented model demonstrates an understanding of some basic "intuitive physics" by registering "surprise" when objects in simulations move in unexpected ways, such as rolling behind a wall and not reappearing on the other side. Credit: Christine Daniloff, MIT

Humans have an early understanding of the laws of physical reality.



Infants, for instance, hold expectations for how objects should move and interact with each other, and will show surprise when they do something unexpected, such as disappearing in a sleight-of-hand magic trick.

Now MIT researchers have designed a <u>model</u> that demonstrates an understanding of some basic "intuitive physics" about how objects should behave. The model could be used to help build smarter artificial intelligence and, in turn, provide information to help scientists understand infant cognition.

The model, called ADEPT, observes objects moving around a scene and makes predictions about how the objects should behave, based on their underlying physics. While tracking the objects, the model outputs a signal at each video frame that correlates to a level of "surprise"—the bigger the signal, the greater the surprise. If an object ever dramatically mismatches the model's predictions—by, say, vanishing or teleporting across a scene—its surprise levels will spike.

In response to videos showing objects moving in physically plausible and implausible ways, the model registered levels of surprise that matched levels reported by humans who had watched the same videos.

"By the time infants are 3 months old, they have some notion that objects don't wink in and out of existence, and can't move through each other or teleport," says first author Kevin A. Smith, a research scientist in the Department of Brain and Cognitive Sciences (BCS) and a member of the Center for Brains, Minds, and Machines (CBMM). "We wanted to capture and formalize that knowledge to build <u>infant cognition</u> into <u>artificial-intelligence</u> agents. We're now getting near human-like in the way models can pick apart basic implausible or plausible scenes."

Joining Smith on the paper are co-first authors Lingjie Mei, an undergraduate in the Department of Electrical Engineering and



Computer Science, and BCS research scientist Shunyu Yao; Jiajun Wu Ph.D. '19; CBMM investigator Elizabeth Spelke; Joshua B. Tenenbaum, a professor of computational <u>cognitive science</u>, and researcher in CBMM, BCS, and the Computer Science and Artificial Intelligence Laboratory (CSAIL); and CBMM investigator Tomer D. Ullman Ph.D. '15.

Mismatched realities

ADEPT relies on two modules: an "inverse graphics" module that captures object representations from raw images, and a "physics engine" that predicts the objects' future representations from a distribution of possibilities.

Inverse graphics basically extracts information of objects—such as shape, pose, and velocity—from pixel inputs. This module captures frames of video as images and uses inverse graphics to extract this information from objects in the scene. But it doesn't get bogged down in the details. ADEPT requires only some approximate geometry of each shape to function. In part, this helps the model generalize predictions to new objects, not just those it's trained on.

"It doesn't matter if an object is rectangle or circle, or if it's a truck or a duck. ADEPT just sees there's an object with some position, moving in a certain way, to make predictions," Smith says. "Similarly, young infants also don't seem to care much about some properties like shape when making physical predictions."

These coarse object descriptions are fed into a physics engine—software that simulates behavior of physical systems, such as rigid or fluidic bodies, and is commonly used for films, video games, and computer graphics. The researchers' physics engine "pushes the objects forward in time," Ullman says. This creates a range of predictions, or a "belief



distribution," for what will happen to those objects in the next frame.

Next, the model observes the actual next frame. Once again, it captures the object representations, which it then aligns to one of the predicted object representations from its belief distribution. If the object obeyed the laws of physics, there won't be much mismatch between the two representations. On the other hand, if the object did something implausible—say, it vanished from behind a wall—there will be a major mismatch.

ADEPT then resamples from its belief distribution and notes a very low probability that the object had simply vanished. If there's a low enough probability, the model registers great "surprise" as a signal spike. Basically, surprise is inversely proportional to the probability of an event occurring. If the probability is very low, the signal spike is very high.

"If an object goes behind a wall, your physics engine maintains a belief that the object is still behind the wall. If the wall goes down, and nothing is there, there's a mismatch," Ullman says. "Then, the model says, 'There's an object in my prediction, but I see nothing. The only explanation is that it disappeared, so that's surprising.'"

Violation of expectations

In development psychology, researchers run "violation of expectations" tests in which infants are shown pairs of videos. One video shows a plausible event, with objects adhering to their expected notions of how the world works. The other video is the same in every way, except objects behave in a way that violates expectations in some way. Researchers will often use these tests to measure how long the infant looks at a scene after an implausible action has occurred. The longer they stare, researchers hypothesize, the more they may be surprised or interested in what just happened.



For their experiments, the researchers created several scenarios based on classical developmental research to examine the model's core object knowledge. They employed 60 adults to watch 64 videos of known physically plausible and physically implausible scenarios. Objects, for instance, will move behind a wall and, when the wall drops, they'll still be there or they'll be gone. The participants rated their surprise at various moments on an increasing scale of 0 to 100. Then, the researchers showed the same videos to the model. Specifically, the scenarios examined the model's ability to capture notions of permanence (objects do not appear or disappear for no reason), continuity (objects move along connected trajectories), and solidity (objects cannot move through one another).

ADEPT matched humans particularly well on videos where objects moved behind walls and disappeared when the wall was removed. Interestingly, the model also matched surprise levels on videos that humans weren't surprised by but maybe should have been. For example, in a video where an object moving at a certain speed disappears behind a wall and immediately comes out the other side, the object might have sped up dramatically when it went behind the wall or it might have teleported to the other side. In general, humans and ADEPT were both less certain about whether that event was or wasn't surprising. The researchers also found traditional neural networks that learn physics from observations—but don't explicitly represent objects—are far less accurate at differentiating surprising from unsurprising scenes, and their picks for surprising scenes don't often align with humans.

Next, the researchers plan to delve further into how infants observe and learn about the world, with aims of incorporating any new findings into their model. Studies, for example, show that infants up until a certain age actually aren't very surprised when objects completely change in some ways—such as if a truck disappears behind a wall, but reemerges as a duck.



"We want to see what else needs to be built in to understand the world more like infants, and formalize what we know about psychology to build better AI agents," Smith says.

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Helping machines perceive some laws of physics (2019, December 2) retrieved 3 May 2024 from <u>https://techxplore.com/news/2019-12-machines-laws-physics.html</u>

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.