


Researchers create a tool for accurately simulating complex systems

May 4 2023, by Adam Zewe



A new technique eliminates a source of bias in a popular simulation method, which could enable scientists to create new algorithms that are more accurate and boost the performance of applications and networks. Credit: Jose-Luis Olivares/MIT

Researchers often use simulations when designing new algorithms, since

testing ideas in the real world can be both costly and risky. But since it's impossible to capture every detail of a complex system in a simulation, they typically collect a small amount of real data that they replay while simulating the components they want to study.

Known as trace-driven simulation (the small pieces of real data are called traces), this method sometimes results in biased outcomes. This means researchers might unknowingly choose an [algorithm](#) that is not the best one they evaluated, and which will perform worse on real data than the simulation predicted that it should.

MIT researchers have developed a new method that eliminates this source of bias in trace-driven simulation. By enabling unbiased trace-driven simulations, the new technique could help researchers design better algorithms for a variety of applications, including improving video quality on the internet and increasing the performance of data processing systems.

The researchers' machine-learning algorithm draws on the principles of causality to learn how the data traces were affected by the behavior of the system. In this way, they can replay the correct, unbiased version of the trace during the simulation.

When compared to a previously developed trace-driven simulator, the researchers' simulation method correctly predicted which newly designed algorithm would be best for video streaming—meaning the one that led to less rebuffering and higher visual quality. Existing simulators that do not account for bias would have pointed researchers to a worse-performing algorithm.

"Data are not the only thing that matter. The story behind how the data are generated and collected is also important. If you want to answer a counterfactual question, you need to know the underlying data

generation story so you only intervene on those things that you really want to simulate," says Arash Nasr-Esfahany, an [electrical engineering](#) and [computer science](#) (EECS) graduate student and co-lead author of a paper on this new technique.

He is joined on the paper by co-lead authors and fellow EECS graduate students Abdullah Alomar and Pouya Hamadianian; recent graduate student Anish Agarwal Ph.D. '21; and senior authors Mohammad Alizadeh, an associate professor of electrical engineering and computer science; and Devavrat Shah, the Andrew and Erna Viterbi Professor in EECS and a member of the Institute for Data, Systems, and Society and of the Laboratory for Information and Decision Systems. The research was recently presented at the [USENIX Symposium on Networked Systems Design and Implementation](#).

Specious simulations

The MIT researchers studied trace-driven simulation in the context of video streaming applications.

In video streaming, an adaptive bitrate algorithm continually decides the video quality, or bitrate, to transfer to a device based on real-time data on the user's bandwidth. To test how different adaptive bitrate algorithms impact network performance, researchers can collect real data from users during a video stream for a trace-driven simulation.

They use these traces to simulate what would have happened to network performance had the platform used a different adaptive bitrate algorithm in the same underlying conditions.

Researchers have traditionally assumed that trace data are exogenous, meaning they aren't affected by factors that are changed during the simulation. They would assume that, during the period when they

collected the network performance data, the choices the bitrate adaptation algorithm made did not affect those data.

But this is often a false assumption that results in biases about the behavior of new algorithms, making the simulation invalid, Alizadeh explains.

"We recognized, and others have recognized, that this way of doing simulation can induce errors. But I don't think people necessarily knew how significant those errors could be," he says.

To develop a solution, Alizadeh and his collaborators framed the issue as a causal inference problem. To collect an unbiased trace, one must understand the different causes that affect the observed data. Some causes are intrinsic to a system, while others are affected by the actions being taken.

In the video streaming example, network performance is affected by the choices the bitrate adaptation algorithm made—but it's also affected by intrinsic elements, like network capacity.

"Our task is to disentangle these two effects, to try to understand what aspects of the behavior we are seeing are intrinsic to the system and how much of what we are observing is based on the actions that were taken. If we can disentangle these two effects, then we can do unbiased simulations," he says.

Learning from data

But researchers often cannot directly observe intrinsic properties. This is where the new tool, called CausalSim, comes in. The algorithm can learn the underlying characteristics of a system using only the trace data.

CausalSim takes trace data that were collected through a randomized control trial, and estimates the underlying functions that produced those data. The model tells the researchers, under the exact same underlying conditions that a user experienced, how a new algorithm would change the outcome.

Using a typical trace-driven simulator, bias might lead a researcher to select a worse-performing algorithm, even though the simulation indicates it should be better. CausalSim helps researchers select the best algorithm that was tested.

The MIT researchers observed this in practice. When they used CausalSim to design an improved bitrate adaptation algorithm, it led them to select a new variant that had a stall rate that was nearly 1.4 times lower than a well-accepted competing algorithm, while achieving the same [video quality](#). The stall rate is the amount of time a user spent rebuffering the video.

By contrast, an expert-designed trace-driven simulator predicted the opposite. It indicated that this new variant should cause a stall rate that was nearly 1.3 times higher. The researchers tested the algorithm on real-world video streaming and confirmed that CausalSim was correct.

"The gains we were getting in the new variant were very close to CausalSim's prediction, while the expert simulator was way off. This is really exciting because this expert-designed simulator has been used in research for the past decade. If CausalSim can so clearly be better than this, who knows what we can do with it?" says Hamadani.

During a 10-month experiment, CausalSim consistently improved [simulation](#) accuracy, resulting in algorithms that made about half as many errors as those designed using baseline methods.

In the future, the researchers want to apply CausalSim to situations where randomized control trial data are not available or where it is especially difficult to recover the causal dynamics of the system. They also want to explore how to design and monitor systems to make them more amenable to causal analysis.

More information: CausalSim: A Causal Framework for Unbiased Trace-Driven Simulation:

www.usenix.org/system/files/nsdi23-alomar.pdf

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: Researchers create a tool for accurately simulating complex systems (2023, May 4) retrieved 6 May 2024 from

<https://techxplore.com/news/2023-05-tool-accurately-simulating-complex.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>
--