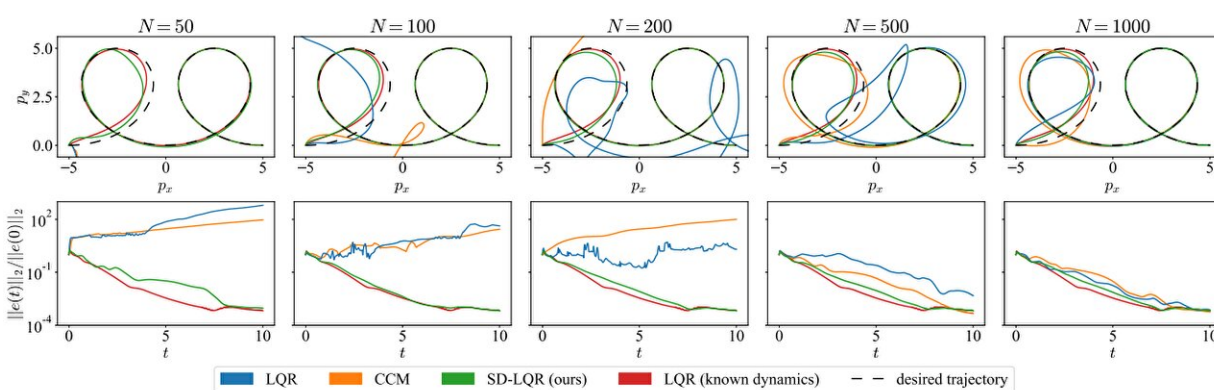


Researchers develop machine-learning technique that can efficiently learn to control a robot

July 26 2023, by Adam Zewe



Trajectory tracking results for the PVTOL system on a double loop-the-loop trajectory. The top row qualitatively depicts the closed-loop trajectories for each method overlaid with the desired trajectory (black dashed). The bottom row shows the normalized tracking error over time. Plots proceed from left to right with an increasing amount N of labeled training data. Our learned SD-LQR method is the only learning-based approach that successfully tracks the trajectory for all N . Credit: *arXiv* (2023). DOI: 10.48550/arxiv.2302.02529

Researchers from MIT and Stanford University have devised a new machine-learning approach that could be used to control a robot, such as a drone or autonomous vehicle, more effectively and efficiently in dynamic environments where conditions can change rapidly.

This technique could help an autonomous vehicle learn to compensate for slippery road conditions to avoid going into a skid, allow a robotic free-flyer to tow different objects in space, or enable a drone to closely follow a downhill skier despite being buffeted by strong winds.

The researchers' approach incorporates certain structure from [control theory](#) into the process for learning a model in such a way that leads to an effective method of controlling complex dynamics, such as those caused by impacts of wind on the trajectory of a flying vehicle. One way to think about this structure is as a hint that can help guide how to control a system.

"The focus of our work is to learn intrinsic structure in the dynamics of the system that can be leveraged to design more effective, stabilizing controllers," says Navid Azizan, the Esther and Harold E. Edgerton Assistant Professor in the MIT Department of Mechanical Engineering and the Institute for Data, Systems, and Society (IDSS), and a member of the Laboratory for Information and Decision Systems (LIDS). "By jointly learning the system's dynamics and these unique control-oriented structures from data, we're able to naturally create controllers that function much more effectively in the real world."

Using this structure in a learned model, the researchers' technique immediately extracts an effective [controller](#) from the model, as opposed to other machine-learning methods that require a controller to be derived or learned separately with additional steps. With this structure, their approach is also able to learn an effective controller using fewer data than other approaches. This could help their learning-based control system achieve better performance faster in rapidly changing environments.

"This work tries to strike a balance between identifying structure in your system and just learning a model from data," says lead author Spencer

M. Richards, a graduate student at Stanford University. "Our approach is inspired by how roboticists use physics to derive simpler models for robots. Physical analysis of these models often yields a useful structure for the purposes of control—one that you might miss if you just tried to naively fit a model to data. Instead, we try to identify similarly useful structure from data that indicates how to implement your control logic."

Additional authors of the paper are Jean-Jacques Slotine, professor of mechanical engineering and of brain and cognitive sciences at MIT, and Marco Pavone, associate professor of aeronautics and astronautics at Stanford. The research will be presented at the [International Conference on Machine Learning](#) (ICML) held July 23–29 in Honolulu. A preprint version is available on the *arXiv* server.

Learning a controller

Determining the best way to control a robot to accomplish a given task can be a [difficult problem](#), even when researchers know how to model everything about the system.

A controller is the logic that enables a drone to follow a desired trajectory, for example. This controller would tell the drone how to adjust its rotor forces to compensate for the effect of winds that can knock it off a stable path to reach its goal.

This drone is a dynamical system—a physical system that evolves over time. In this case, its position and velocity change as it flies through the environment. If such a system is simple enough, engineers can derive a controller by hand.

Modeling a system by hand intrinsically captures a certain structure based on the physics of the system. For instance, if a robot were modeled manually using differential equations, these would capture the

relationship between velocity, acceleration, and force. Acceleration is the rate of change in velocity over time, which is determined by the mass of and forces applied to the robot.

But often the system is too complex to be exactly modeled by hand. Aerodynamic effects, like the way swirling wind pushes a flying vehicle, are notoriously difficult to derive manually, Richards explains. Researchers would instead take measurements of the drone's position, velocity, and rotor speeds over time, and use machine learning to fit a model of this dynamical system to the data.

But these approaches typically don't learn a control-based structure. This structure is useful in determining how to best set the rotor speeds to direct the motion of the drone over time.

Once they have modeled the dynamical system, many existing approaches also use data to learn a separate controller for the system.

"Other approaches that try to learn dynamics and a controller from data as separate entities are a bit detached philosophically from the way we normally do it for simpler systems. Our approach is more reminiscent of deriving models by hand from physics and linking that to control," Richards says.

Identifying structure

The team from MIT and Stanford developed a technique that uses machine learning to learn the dynamics model, but in such a way that the model has some prescribed structure that is useful for controlling the system.

With this structure, they can extract a controller directly from the dynamics model, rather than using data to learn an entirely separate

model for the controller.

"We found that beyond learning the dynamics, it's also essential to learn the control-oriented structure that supports effective controller design. Our approach of learning state-dependent coefficient factorizations of the dynamics has outperformed the baselines in terms of data efficiency and tracking capability, proving to be successful in efficiently and effectively controlling the system's trajectory," Azizan says.

When they tested this approach, their controller closely followed desired trajectories, outpacing all the baseline methods. The controller extracted from their learned model nearly matched the performance of a ground-truth controller, which is built using the exact dynamics of the system.

"By making simpler assumptions, we got something that actually worked better than other complicated baseline approaches," Richards adds.

The researchers also found that their method was data-efficient, which means it achieved high performance even with few data. For instance, it could effectively model a highly dynamic rotor-driven vehicle using only 100 data points. Methods that used multiple learned components saw their performance drop much faster with smaller datasets.

This efficiency could make their technique especially useful in situations where a drone or robot needs to learn quickly in rapidly changing conditions.

Plus, their approach is general and could be applied to many types of dynamical systems, from robotic arms to free-flying spacecraft operating in low-gravity environments.

In the future, the researchers are interested in developing models that are more physically interpretable, and that would be able to identify very

specific information about a dynamical system, Richards says. This could lead to better-performing controllers.

"Despite its ubiquity and importance, nonlinear feedback control remains an art, making it especially suitable for data-driven and learning-based methods. This paper makes a significant contribution to this area by proposing a method that jointly learns system dynamics, a controller, and control-oriented structure," says Nikolai Matni, an assistant professor in the Department of Electrical and Systems Engineering at the University of Pennsylvania, who was not involved with this work.

"What I found particularly exciting and compelling was the integration of these components into a joint learning algorithm, such that control-oriented structure acts as an inductive bias in the learning process. The result is a data-efficient learning process that outputs dynamic models that enjoy intrinsic structure that enables effective, stable, and robust control. While the technical contributions of the paper are excellent themselves, it is this conceptual contribution that I view as most exciting and significant," says Matni.

More information: Spencer M. Richards et al, Learning Control-Oriented Dynamical Structure from Data, *arXiv* (2023). [DOI: 10.48550/arxiv.2302.02529](https://doi.org/10.48550/arxiv.2302.02529)

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 develop machine-learning technique that can efficiently learn to control a robot (2023, July 26) retrieved 9 May 2024 from <https://techxplore.com/news/2023-07-machine-robot>

learning-technique-efficiently-robot.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.