

Advanced series of more robust drones are teaching themselves how to fly

March 13 2020, by Caitlin Dawson



A robust drone controller, trained using reinforcement learning algorithms developed by USC researchers allows drones to withstand harsh treatment, including being thrown and pushed. Credit: Haotian Mai

Drones, specifically quadcopters, are an adaptable lot. They've been used to assess damage after disasters, deliver ropes and life-jackets in areas too dangerous for ground-based rescuers, survey buildings on fire and deliver medical specimens.

But to achieve their full potential, they have to be tough. In the [real world](#), drones are forced to navigate uncertain shapes in collapsing buildings, avoid obstacles and deal with challenging conditions, including storms and earthquakes.

At the USC Viterbi School of Engineering's Department of Computer Science, researchers have created artificially intelligent drones that can quickly recover when pushed, kicked or when colliding with an object. The autonomous drone "learns" how to recover from a slew of challenging situations thrown at it during a simulation process.

"Currently, the controllers designed to stabilize quadcopters require careful tuning and even then, they are limited in terms of robustness to disruption and are model-specific," said the study's lead author Artem Molchanov, a Ph.D. in computer science candidate in USC's Robotic Systems Embedded Laboratory.

"We're trying to eliminate this problem and present an approach that leverages recent advancement in reinforcement learning so we can completely eliminate hand-tuning controllers and make drones super robust to disruptions."

The paper, called "Sim-to-(Multi)-Real: Transfer of Low-Level Robust Control Policies to Multiple Quadrotors," was presented at the International Conference on Intelligent Robots and Systems.

Co-authors were Tao Chen, USC computer science master's student; Wolfgang Honig, a former USC computer science Ph.D. student; James A. Preiss, a computer science Ph.D. student; Nora Ayanian, USC assistant professor of computer science and Andrew and Erna Viterbi Early Career Chair; and Gaurav Sukhatme, professor of computer science and electrical and computer engineering and USC Viterbi executive vice dean.

Learning to fly

Roboticians have been turning to birds for flight inspiration for years. But drones have a long way to go before they're as agile as their feathered counterparts. When a drone ends up in an undesirable orientation, such as upside down, it can be difficult for it to right itself. "A drone is an inherently unstable system," said Molchanov.

"Controlling a drone requires a lot of precision. Especially when something sudden occurs, you need a fast and precise sequence of control inputs." But, if a drone was able to learn from experience, like humans, it would be more capable of overcoming these challenges.

With this in mind, the USC researcher team created a system that uses a type of machine learning, a subset of artificial intelligence, called reinforcement learning to train the drone in a simulated environment. More precisely, to train the drone's "brain," or neural network controller.

"Reinforcement learning is inspired by biology—it's very similar to how you might train a dog with a reward when it completes a command," said Molchanov.

Of course, drones don't get snacks. But in the process of [reinforcement learning](#), they do receive an algorithmic reward: a mathematical reinforcement signal, which is positive reinforcement that it uses to infer which actions are most desirable.



Lead author Artem Molchanov (center), a PhD in computer science student, with co-authors Professor Gaurav Sukhatme and Assistant Professor Nora Ayanian. Credit: Haotian Mai

Learning in simulation

The drone starts in simulation mode. At first, it knows nothing about the world or what it is trying to achieve, said Molchanov. It tries to jump a little bit or rotate on the ground.

Eventually, it learns to fly a little bit and receives the positive reinforcement signal. Gradually, through this process, it understands how to balance itself and ultimately fly. Then, things get more complicated.

While still in simulation, the researchers throw randomized conditions at the controller until it learns to deal with them successfully. They add noise to the input to simulate a realistic sensor. They change the size and

strength of the motor and push the drone from different angles.

Over the course of 24 hours, the system processes 250 hours of real-world training. Like training wheels, learning in simulation mode allows the drone to learn on its own in a safe environment, before being released into the wild. Eventually, it finds solutions to every challenge put in its path.

"In simulation we can run hundreds of thousands of scenarios," said Molchanov.

"We keep slightly changing the simulator, which allows the drone to learn to adapt to all possible imperfections of the environment."



Three quadcopters of different sizes controlled by the same policy trained entirely in simulation. Credit: Molchanov et al.

A real-world challenge

To prove their approach, the researchers moved the trained controller onto real drones developed in Ayanian's Automatic Coordination of Teams Lab. In a netted indoor [drone](#) facility, they flew the drones and tried to throw them off by kicking and pushing them.

The drones were successful in correcting themselves from moderate hits (including pushes, light kicks and colliding with an object) 90% of the time. Once trained on one machine, the controller was able to quickly generalize to quadcopters with different dimensions, weights and sizes.

While the researchers focused on robustness in this study, they were surprised to find the system also performed competitively in terms of trajectory tracking—moving from point A to B to C. While not specifically trained for this purpose, it seems the rigorous simulation training also equipped the controller to follow a [moving target](#) precisely.

The researchers note that there's still work to be done. In this experiment, they manually adjusted a few parameters on the drones, for example, limiting maximum thrust, but the next step is to make the drones completely independent. The experiment is a promising move towards building sturdy drones that can tune themselves and learn from experience.

Professor Sukhatme, Molchanov's advisor and a Fletcher Jones Foundation Endowed Chair in Computer Science, said the research solves two important problems in robotics: robustness and generalization.

"From a safety perspective, robustness is super important. If you're building a flight control system, it can't be brittle and fall apart when something goes wrong," said Sukhatme.

"The other important thing is generalization. Sometimes you might build a very safe system, but it will be very specialized. This research shows what a mature and accomplished Ph.D. student can achieve, and I'm very proud of Artem and the team he assembled."

Provided by University of Southern California

Citation: Advanced series of more robust drones are teaching themselves how to fly (2020, March 13) retrieved 18 April 2024 from <https://techxplore.com/news/2020-03-advanced-series-robust-drones.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.