PennState College of Engineering | Systems, Controls, and Robotics Seminar Series

Safe Visual Reinforcement Learning via Conformal Prediction

Osbert Bastani

Department of Computer and Information Science
University of Pennsylvania

Seminar: Friday, February 9
2 PM, 125 Reber

ABSTRACT

Reinforcement learning is a promising approach to solving hard robotics tasks, yet an important obstacle to deploying reinforcement learning is the difficulty in ensuring safety. We build on an approach that composes the learned policy with a backup policy: it uses the learned policy on the interior of the region where the backup policy is guaranteed to be safe, and switches to the backup policy on the boundary of this region. The key challenge is checking when the backup policy is guaranteed to be safe. First, we propose model predictive shielding (MPS), which uses online model-based verification to efficiently perform this check. For visual control, such a model-based approach is difficult to employ; instead, we propose an algorithm that uses deep learning to predict when a state may be unsafe. Then, we use a statistical technique called conformal prediction to provide high-probability safety guarantees for this approach.

BIOGRAPHY

Osbert Bastani is an assistant professor at the Department of Computer and Information Science at the University of Pennsylvania. He is broadly interested in techniques for designing trustworthy machine learning systems. Previously, he completed his Ph.D. in computer science from Stanford and his A.B. in mathematics from Harvard.