

# THE COMPUTER SCIENCE DISTINGUISHED COLLOQUIUM



## Peter Stone

Dr. Peter Stone is the David Bruton, Jr. Centennial Professor and Associate Chair of Computer Science, as well as Director of Texas Robotics, at the University of Texas at Austin. He is also University Distinguished Teaching Professor. Dr. Stone's research interests in AI include machine learning and reinforcement learning, multiagent systems, and robotics. Dr. Stone received his Ph.D in CS in 1998 from Carnegie Mellon University. He was a Senior Technical Staff Member in the Artificial Intelligence Principles Research Department at AT&T Labs - Research. He is an Alfred P. Sloan Research Fellow, Guggenheim Fellow, AAAI Fellow, IEEE Fellow, AAAS Fellow, ACM Fellow, Fulbright Scholar, and 2004 ONR Young Investigator. Dr. Stone co-founded Cogitai, Inc., a startup company focused on continual learning, in 2015, and currently serves as Executive Director of Sony AI America.



**Date: Friday, April 15**

**Time: 12:30 – 1:30 PM**

**Zoom: <https://bit.ly/3JOsLut>**

**(Virtual Event)**

### Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning

Many potential applications of artificial intelligence involve making real-time decisions in physical systems while interacting with humans. Automobile racing represents an extreme example of these conditions; drivers must execute complex tactical manoeuvres to pass or block opponents while operating their vehicles at their traction limits. Racing simulations, such as the PlayStation game Gran Turismo, faithfully reproduce the non-linear control challenges of real race cars while also encapsulating the complex multi-agent interactions. Here we describe how we trained agents for Gran Turismo that can compete with the world's best e-sports drivers. We combine state-of-the-art, model-free, deep reinforcement learning algorithms with mixed-scenario training to learn an integrated control policy that combines exceptional speed with impressive tactics. In addition, we construct a reward function that enables the agent to be competitive while adhering to racing's important, but under-specified, sportsmanship rules. We demonstrate the capabilities of our agent, Gran Turismo Sophy, by winning a head-to-head competition against four of the world's best Gran Turismo drivers. By describing how we trained championship-level racers, we demonstrate the possibilities and challenges of using these techniques to control complex dynamical systems in domains where agents must respect imprecisely defined human norms.