

Instance Dependent Sample Complexity Bounds for Interactive Learning

Dr. Kevin Jamieson

Assistant Professor

Paul G. Allen School of Computer Science & Engineering
University of Washington

Abstract: The sample complexity of an interactive learning problem, such as multi-armed bandits or reinforcement learning, is the number of interactions with nature required to output an answer (e.g., a recommended arm or policy) that is approximately close to optimal with high probability. While minimax guarantees can be useful rules of thumb to gauge the difficulty of a problem class, algorithms optimized for this worst-case metric often fail to adapt to “easy” instances where fewer samples suffice. In this talk, I will highlight some of my group’s work on algorithms that obtain optimal, finite time, instance dependent sample complexities that scale with the true difficulty of the particular instance, versus just the worst-case. In particular, I will describe a unifying experimental design based approach used to obtain such algorithms for best-arm identification for linear bandits, contextual bandits with arbitrary policy classes, and smooth losses for linear dynamical systems.

Bio: Kevin Jamieson is an Assistant Professor in the Paul G. Allen School of Computer Science & Engineering at the University of Washington and is the Guestrin Endowed Professor in Artificial Intelligence and Machine Learning. He received his B.S. in 2009 from the University of Washington, his M.S. in 2010 from Columbia University, and his Ph.D. in 2015 from the University of Wisconsin - Madison under the advisement of Robert Nowak, all in electrical engineering. He returned to the University of Washington as faculty in 2017 after a postdoc with Benjamin Recht at the University of California, Berkeley. Jamieson’s research explores how to leverage already-collected data to inform what future measurements to make next, in a closed loop. His work ranges from theory to practical algorithms with guarantees to open-source machine learning systems and has been adopted in a range of applications, including measuring human perception in psychology studies, adaptive A/B/n testing in dynamic web-environments, numerical optimization, and efficient tuning of hyperparameters for deep neural networks.