



Yaqi Duan

**Department of Technology, Operations and Statistics
New York University**

Optimal policy evaluation using kernel-based temporal difference methods

Wednesday, October 30th, 2024

12:00 PM

110 Frelinghuysen Road, Hill Center, Room 116

Zoom Meeting: Meeting ID: 969 0606 4706

Password: 745339

<https://rutgers.zoom.us/j/96906064706?pwd=ZklvbExpRVBJQ3c5dUhhYTFuR2ZrZz09>

Light refreshments will be served in Hill 452, 11:20am

Abstract:

In this talk, we study methods based on reproducing kernel Hilbert spaces (RKHS) for estimating the value function of an infinite-horizon discounted Markov reward process (MRP). We analyze the error of this estimate in the $L^2(\mu)$ -norm, where μ denotes the stationary distribution of the underlying Markov chain. We use empirical process theory techniques to derive a non-asymptotic upper bound on the error with explicit dependence on the eigenvalues of the associated kernel operator, as well as the instance-dependent variance of the Bellman residual error. In addition, we establish minimax lower bounds, showing that our rates are optimal in terms of sample size n and the effective horizon $H = (1 - \gamma)^{-1}$. While existing worst-case results predict cubic scaling in H , our analysis reveals a wide range of possible scalings, with only parametric cases reaching the worst-case cubic rate.

Our analysis relies on a critical inequality that captures the trade-off between variance and bias, while accounting for the dynamic structure of the system. The theory reveals an interplay between the curse of horizon and the curse of dimensionality, providing new insights into policy evaluation in reinforcement learning.

Bio: Yaqi Duan is an Assistant Professor at the Department of Technology, Operations, and Statistics at Stern School of Business at New York University. Duan's research interests lie at the intersection of statistics, machine learning and operations research, with a focus on statistical methodologies and theories to address challenges in data-driven decision-making problems. Her work in reinforcement learning was honored with the 2023 IMS Lawrence D. Brown Ph.D. Student Award. Duan earned her Ph.D. from Princeton University and a B.S. in Mathematics from Peking University. Prior to joining NYU Stern, Duan was a postdoctoral researcher at MIT, working with Professor Martin J. Wainwright.

