Multi-Dimensional Distributional Reinforcement Learning: A Hilbert Space Embedding Approach

Mehrdad Mohammadi[1], **Qi Zheng[2]**, Ruoqing Zhu[3]

[1]*Department of Statistics, University of Illinois Urbana-Champaign*

[2]*Department of Bioinformatics and Biostatistics, University of Louisville*

[3]*Department of Statistics, University of Illinois Urbana-Champaign*

ABSTRACT

We propose an (offline) distributional reinforcement learning framework (RK-DRL) that leverages Hilbert space embeddings to estimate the multi-dimensional value distribution under a proposed target policy. In our setting the state-action are also multi-dimensional and continuous. By mapping probability measures into a reproducing kernel Hilbert space via kernel mean embeddings, our method replaces Wasserstein metrics with a novel integral probability metric. This enables efficient estimation in multi-dimensional state–action spaces and reward settings, where direct computation of Wasserstein distances is computationally challenging. Theoretically, we establish contraction properties of the distributional Bellman operator under our proposed metric involving the Mat´ern family of kernels and provide uniform con-vergence guarantees. Empirical results demonstrate improved convergence rates and robust off-policy evaluation under mild assumptions, namely, Lipschitz continuity and boundedness for the kernels, highlighting the potential of our embedding-based approaches in complex, real-world decision-making scenarios and risk evaluations.

Keywords: Wasserstein Distance, Reproducing Kernel Hilbert Space, Non-parametric, Matern Kernel

Tuning Parameter Calibration for Prediction in Personalized Medicine

Shih-Ting Huang

Department of Bioinformatics and Biostatistics, University of Louisville

ABSTRACT

Personalized medicine has become an important part of medicine, for instance predicting individual drug responses based on genomic information. However, many current statistical methods are not tailored to this task, because they overlook the individual heterogeneity of patients. In this paper, we look at personalized medicine from a linear regression standpoint. We introduce an alternative version of the ridge estimator and target individuals by establishing a tuning parameter calibration scheme that minimizes prediction errors of individual patients. In stark contrast, classical schemes such as cross-validation minimize prediction errors only on average. We show that our pipeline is optimal in terms of oracle inequalities, fast, and highly effective both in simulations and on real data.

Keywords: Personalized Prediction; Ridge Regression; Precision Medicine; Parameter Calibration