An Integrated GMM Shrinkage Approach with Consistent Moment Selection from Multiple External Sources

Jun Shao

University of Wisconsin-Madison

ABSTRACT

Interest has grown in analyzing primary internal data by utilizing some independent external aggregated statistics for efficiency gain. However, when population heterogeneity exists, inappropriate incorporation may lead to a biased estimator. With multiple external sources under generalized estimation equations and possibly heterogeneous populations, we propose an integrated generalized moment method that can perform a data-driven selection of valid moment equations from external sources and make efficient parameter estimation simultaneously. Moment equation selection consistency and asymptotic normality are established for the proposed estimator. Further, when the sample sizes of all external sources are large compared to the internal sample size, asymptotically the proposed estimator is more efficient than the estimator based on the internal data only and is oracle-efficient in the sense that it is as efficient as the oracle estimator based on all valid moment equations. Simulation studies confirm the theoretical results and the efficiency of the proposed method empirically. An example is also included for illustration.

Keywords: Adaptive lasso; Data integration; Generalized method of moments; Heterogeneous population

Training-Free Multi-Agent Language Models

Xiaowu Dai

Departments of Statistics and Data Science and of Biostatistics, UCLA

ABSTRACT

Large Language Models (LLMs) have demonstrated strong generative capabilities but remain prone to inconsistencies and hallucinations. We introduce Peer Elicitation Games (PEG), a training-free, game-theoretic framework for aligning LLMs through a peer elicitation mechanism involving a generator and multiple discriminators instantiated from distinct base models. Discriminators interact in a peer evaluation setting, where rewards are computed using a determinant-based mutual information score that provably incentivizes truthful reporting without requiring ground-truth labels. We establish theoretical guarantees showing that each agent, via online learning, achieves sublinear regret in the sense their cumulative performance approaches that of the best fixed truthful strategy in hindsight. Moreover, we prove last-iterate convergence to a truthful Nash equilibrium, ensuring that the actual policies used by agents converge to stable and truthful behavior over time. Empirical evaluations across multiple benchmarks demonstrate significant improvements in factual accuracy. These results position PEG as a practical approach for eliciting truthful behavior from LLMs without supervision or fine-tuning.

Keywords: LLMs; Nash equilibrium; No-regret learning; Incentives

Optimal-PhiBE for Continuous-Time Reinforcement Learning with Discrete-Time Data

Yuhua Zhu[1]

Department of Statistics and Data Science, University of California-Los Angeles

ABSTRACT

This talk addresses continuous-time reinforcement learning (RL) in settings where the system dynamics are governed by a stochastic differential equation but remains unknown, with only discrete-time observations available. We introduce Optimal-PhiBE, an equation that integrates discrete-time information into a PDE, combining the strengths of both RL and PDE formulations. In linear-quadratic control, Optimal-PhiBE can even achieve accurate continuous-time optimal policy with only discrete-time information. In general dynamics, Optimal-PhiBE is less sensitive to reward oscillations, leading to smaller discretization errors. Furthermore, we extend Optimal-PhiBE to higher orders, providing increasingly accurate approximations. At the end of the talk, I will discuss how this technique can be leveraged to generate time-dependent samples and tackle goal-oriented inverse problems.

Keywords: Reinforcement learning; Optimal control; Time-series data