For favour of Posting

THE HONG KONG UNIVERSITY OF SCIENCE & TECHNOLOGY

Department of Mathematics

SEMINAR ON DATA SCIENCE AND APPLIED MATHEMATICS

Advancements in Kernel Learning and Offline Reinforcement Learning through Generative Models

By

Prof. Wenjia WANG
HKUST-GZ

Abstract

In this talk for MATH 5470, I will present two classes of my recent research.

In Part I, I will talk about random smoothing data augmentation. Random smoothing data augmentation is a unique form of regularization that can prevent overfitting by introducing noise to the input data, encouraging the model to learn more generalized features. In this work, we aim to bridge this gap by presenting a framework for random smoothing regularization that can adaptively and effectively learn a wide range of ground truth functions belonging to the classical Sobolev spaces. By using random smoothing regularization as novel convolution-based smoothing kernels, we can attain optimal convergence rates in these cases using a kernel gradient descent algorithm, either with early stopping or weight decay.

In Part II, I will talk about our recent series of works on offline reinforcement learning.
Due to the inability to interact with the environment, offline reinforcement learning (RL) methods face the challenge of estimating the Out-of-Distribution (OOD) points. Existing methods for addressing this issue either control policy to exclude the OOD action or make the Q-function pessimistic. However, these methods can be overly conservative or fail to identify OOD areas accurately. In this talk, I will be discussing our recent advancements in offline reinforcement learning, specifically focusing on the utilization of generative models such as GAN and diffusion models. Our proposed methods are evaluated on the D4RL benchmarks and have demonstrated significant improvements across numerous tasks. Theoretical results are provided for performance guarantee.

References:
1. Ding, L., Hu, T., Jiang, J., Li, D., Wang, W., & Yao, Y. (2024). Random smoothing regularization in kernel gradient descent learning. Journal of Machine Learning Research.
2. Fang, L., Liu, R., Zhang, J., Wang, W., & Jing, B. Y. (2025). Diffusion Actor-Critic: Formulating Constrained Policy Iteration as Diffusion Noise Regression for Offline Reinforcement Learning. The Thirteenth International Conference on Learning Representations (ICLR).
3. Zhang, J., Fang, L., Shi, K., Wang, W., & Jing, B. Y. (2024). Q-Distribution guided Q-learning for offline reinforcement learning: Uncertainty penalized Q-value via consistency model. Neural Information Processing Systems (NeurIPS), 2024.
4. Zhang, J., Zhang, C., Wang, W., & Jing, B. Y. (2023). Constrained Policy Optimization with Explicit Behavior Density For Offline Reinforcement Learning. Neural Information Processing Systems (NeurIPS), 2023.

Bio: Wenjia Wang is an assistant professor in the Data Science and Analysis Thrust at the Information Hub of the Hong Kong University of Science and Technology (Guangzhou). He obtained his Ph.D. in the School of Industrial & Systems Engineering at Georgia Institute of Technology. Wenjia Wang's research interests include uncertainty quantification, computer experiments, machine learning, stochastic simulation, and nonparametric statistics.

Date : 7 April 2025 (Monday)
Time : 7:00p.m. (HKTime)

Join Zoom Meeting:
https://hkust.zoom.us/j/98329223175?pwd=di3nMuDReJjsFaoyrddvE1OjDotBoj.1
Meeting ID: 983 2922 3175
Passcode: hkust