

The Hong Kong University of Science and Technology

Seminar on Business Data Science

Department of ISOM

Exploring the Training Robustness of Distributional Reinforcement Learning against Noisy State Observations

By

Prof. Linglong KONG

University of Alberta

Date: 15 January, 2024 (Monday)

Time: 2:00pm – 3:00pm

Venue: Meeting Room 4047 (LSK Business Building)

Abstract

In real scenarios, state observations that an agent observes may contain measurement errors or adversarial noises, misleading the agent to take suboptimal actions or even collapse while training. In this paper, we study the training robustness of distributional Reinforcement Learning (RL), a class of state-of-the-art methods that estimate the whole distribution, as opposed to only the expectation, of the total return. Firstly, we validate the contraction of distributional Bellman operators in the State-Noisy Markov Decision Process (SN-MDP), a typical tabular case that incorporates both random and adversarial state observation noises. In the noisy setting with function approximation, we then analyze the vulnerability of least squared loss in expectation-based RL with either linear or nonlinear function approximation. By contrast, we theoretically characterize the bounded gradient norm of distributional RL loss based on the categorical parameterization equipped with the Kullback–Leibler (KL) divergence. The resulting stable gradients while the optimization in distributional RL accounts for its better training robustness against state observation noises. Finally, extensive experiments on the suite of environments verified that distributional RL is less vulnerable against both random and adversarial noisy state observations compared with its expectation-based counterpart.

Bio

Dr. Linglong Kong is a professor in the Department of Mathematical and Statistical Sciences at the University of Alberta. He holds a Canada Research Chair in Statistical Learning, a Canada CIFAR AI Chair, and is a fellow of the Alberta Machine Intelligence Institute (AMII). His publication record includes about 100 peer-reviewed articles in top journals such as AOS, JASA and JRSSB as well as top conferences such as NeurIPS, ICML, ICDM, AAAI, and IJCAI. Dr. Kong currently serves as associate editor of the Journal of the American Statistical Association, the Canadian Journal of Statistics, and Statistics and its Interface, as well as guest editor of Statistics and its Interface. Additionally, Dr. Kong is a member of the Regional Advisory Board of the Western North American Region of the International Biometric Society. He served as a guest editor of Canadian Journal of Statistics, associate editor of International Journal of Imaging Systems and Technology, guest associate editor of Frontiers of Neurosciences, chair of the ASA Statistical Imaging Session, member of the Statistics Society of Canada's Board of Directors, chair of the ASA Statistical Computing Session program, and chair of the webinar committee. He is interested in the analysis of high-dimensional and neuroimaging data, statistical machine learning, robust statistics and quantile regression, as well as artificial intelligence for smart health.

All interested are welcome!

Enquiries: Dept of ISOM