

Learning to Optimally Stop Diffusion Processes, with Financial Applications

We study optimal stopping for diffusion processes with unknown model primitives within the continuous-time reinforcement learning (RL) framework developed by Wang et al. (2020), and present applications to option pricing and portfolio choice. By penalizing the corresponding variational inequality formulation, we transform the stopping problem into a stochastic optimal control problem with two actions. We then randomize controls into Bernoulli distributions and add an entropy regularizer to encourage exploration. We derive a semi-analytical optimal Bernoulli distribution, based on which we devise RL algorithms using the martingale approach established in Jia and Zhou (2022). We establish a policy improvement theorem and prove the fast convergence of the resulting policy iterations. We demonstrate the effectiveness of the algorithms in pricing finite-horizon American put options, solving Merton's problem with transaction costs, and scaling to high-dimensional optimal stopping problems. In particular, we show that both the offline and online algorithms achieve high accuracy in learning the value functions and characterizing the associated free boundaries. Joint work with Min Dai, Yu Sun and Zuoquan Xu.

22 August, 2025 (Fri) 10:00am



SEMINAR LINK

cityu.zoom.us/j/86217436061
Meeting ID: 862 1743 6061

Speaker

Professor Xunyu ZHOU

Department of Industrial Engineering
and Operations Research (IEOR)
Columbia University, USA

- • • Xunyu Zhou is the Liu Family Professor of Industrial Engineering and Operations Research at Columbia University in New York. Before joining Columbia, he was the Nomura Professor of Mathematical Finance, the Director of Nomura Center for Mathematical Finance and the Director of Oxford-Nie Financial Big Data Lab at University of Oxford during 2007-2016, and Choh-Ming Li Professor of Systems Engineering and Engineering Management at The Chinese University of Hong Kong during 2013-2014.

He is well known for his work in indefinite stochastic LQ control theory and application to dynamic mean–variance portfolio selection, in asset allocation and pricing under cumulative prospect theory, and in general time inconsistent problems. His current research focuses on reinforcement learning for controlled diffusion processes and applications to generative AI and intelligent wealth management solutions. He directs the Nie Center for Intelligent Asset Management, a research center funded by a FinTech company, at Columbia. He has addressed the 2010 International Congress of Mathematicians, and has been awarded the Wolfson Research Award from The Royal Society (UK), the Outstanding Paper Prize from the Society for Industrial and Applied Mathematics, the Humboldt Distinguished Lecturer, the Alexander von Humboldt Research Fellowship, the Archimedes Lecturer at Columbia, and Distinguished Faculty Teaching Award at Columbia University. He is both an IEEE Fellow and a SIAM Fellow.

Professor Zhou received his Ph.D. in Operations Research and Control Theory from Fudan University in China in 1989.

Enquiry : 3442 8422 | All are welcome



Seminar 2024–2025, SYE32