香港浸會大學
HONG KONG BAPTIST UNIVERSITY

DEPARTMENT OF
COMPUTER SCIENCE
計算機科學系

TMLR Young Scientist SEMINAR

2026 SERIES

Trustworthy Machine Learning and Reasoning Group

Dr. Ziyang Luo

Research Scientist,
Salesforce AI Research

Date: 11 May 2026 (Monday)

Time: 13:30 – 15:00 (HKT)

Meeting: https://hkbu.zoom.us/j/6603117755

Agentic RL: Lessons from Building the Full Training Loop

ABSTRACT

The hard part of agentic RL is rarely a specific loss function; it is keeping the environment, rollout, and training loop usefully coupled as models move toward long-horizon and tool-heavy interaction. This talk shares hands-on lessons from building the full training loop, organized around three parts.

(1) *Environment*. Why structural fidelity, verifier design, and scaffold diversity, rather than dataset size, determine whether RL produces any learnable signal, and what "environment as the first scaling axis" means in practice.

(2) *Rollout*. How long and uneven trajectories involving tool calls reshape sampling, scheduling, and context management, and why partial rollout, stale sample filtering, and context transitions stop being infrastructure plumbing and become part of the algorithm.

(3) *Fully async. vs sync. training*. The real trilemma among throughput, stability, and agent flexibility, and how windowed scheduling, TITO consistency, and alignment between training and serving decide whether the learner is optimizing the distribution you actually deploy.

BIOGRAPHY

Ziyang Luo is a Research Scientist at Salesforce AI Research, working on agentic reinforcement learning. He received his PhD from Hong Kong Baptist University and was a visiting researcher at the National University of Singapore. He previously interned at Microsoft Research and Alibaba. Homepage: https://chiyeunglaw.github.io/

ENQUIRY

Email: bhanml@comp.hkbu.edu.hk