

Low, Siow Meng

ML researcher — safe RL with feedback; exploring RL × LLMs

Singapore • GMT+8 • Singapore Citizen

✉ siowmeng@siowmenglow.com

🌐 github.com/siowmeng • [linkedin.com/in/siowmeng](https://www.linkedin.com/in/siowmeng) • [🎓 Google Scholar](https://scholar.google.com/citations?user=siowmeng)

☎ +65 9233 1171

Research Summary

- Focus: safe reinforcement learning and alignment from **sparse / weak feedback** (e.g., pass/fail safety labels), with interest in RL×LLMs.
- **TraCeS** turns rollout-level safety labels into per-step safety signals, enabling agents to learn safer behavior without hand-designed safety costs.
- Developed methods for safe learning/planning under implicit or hard-to-specify constraints (AAAI'22, ICAPS'23, arXiv'24; ICLR'25 workshop).

Publications

Under review / working paper

- **Low, S. M.**, Gong, Z., Kumar, A. *TraCeS: Learning per-timestep constraint-violation credit from sparse trajectory-level labels* (extended version; under review).

Refereed and public manuscripts

- **Low, S. M.**, Kumar, A. *TraCeS: Trajectory Based Credit Assignment From Sparse Safety Feedback*. ICLR 2025 Workshop on Bidirectional Human-AI Alignment (non-archival). [\[PDF\]](#)
- **Low, S. M.**, Kumar, A. *Safe Reinforcement Learning with Learned Non-Markovian Safety Constraints*. arXiv:2405.03005. [\[PDF\]](#)
- **Low, S. M.**, Kumar, A., Sanner, S. *Safe MDP planning by learning temporal patterns of undesirable trajectories and averting negative side effects*. ICAPS 2023. [\[PDF\]](#)
- **Low, S. M.**, Kumar, A., Sanner, S. *Sample-efficient iterative lower bound optimization of deep reactive policies for planning in continuous MDPs*. AAAI 2022. [\[Publisher\]](#)
- Xiao, W., **Low, S. M.**, Tham, C. K., Das, S. *Prediction based energy-efficient task allocation for delay-constrained wireless sensor networks*. SECON Workshops 2009. [\[Publisher\]](#)

Education

Singapore Management University (SMU), Singapore — PhD, Computer Science 2020 – 2025

Thesis: *From sparse feedback to sequential decision-making: Learning safety constraints with weak supervision*. ([thesis link](#))

Awards: SMU Presidential Doctoral Fellowship in Computing (AY2022/23, AY2023/24).

Imperial College London, United Kingdom — MSc, Business Analytics 2016 – 2017

Distinction; Dean's List; Book Prize (Best overall performance).

National University of Singapore (NUS), Singapore — BEng, Electrical Engineering 2004 – 2008

First Class Honours; NUS Centre of Intelligent Control Book Prize (2006); Dean's List (2005, 2006).

Experience

Singapore Management University (SMU), Singapore — Adjunct Lecturer Aug 2025 – Present

- COR-IS1704 Computational Thinking and Programming (Python), AY2024/25 T1. Final rating: above 90th percentile across SMU faculty.
- CS205 Operating Systems (undergraduate core), AY2025/26 T1.

- IS112 Data Management (ER modeling, relational schema, SQL, indexing / performance), AY2025/26 T2.

Intech Process Automation, Singapore — IIoT Data Science Lead *2019 – 2020*

- Led applied ML work for industrial predictive analytics; designed product UI and built proof-of-concept workflows for data prep and model building.

Amaris.AI, Singapore — AI Data Scientist *2018 – 2019*

- Built real-time social media text mining + sentiment platform; delivered automated forecasting solution for an electricity market operator.

NEC Laboratories Singapore — Researcher (Intelligent Data Analytics) *2017 – 2018*

- Built ML models for risk profiling; improved prediction accuracy by 50% over prior heuristics; delivered interpretable factor visualizations for intervention.

IBM, Singapore — Software Presales Consultant *2013 – 2015*

Selected outcomes: designed proofs-of-concept; contributed to annual software revenue ~US\$1M.

Deloitte, Singapore — Senior Software Consultant *2008 – 2013*

Led delivery teams; designed/integrated solutions across 30+ systems; improved productivity metrics by ~6×.

Service

Reviewer: NeurIPS (2024–25); ICLR (2025–26); ICML (2025); AAAI (2023–26); ICAPS Planning & Learning Track (2022–23); AISTATS (2025); JAIR (2023–25).

Singapore Computer Society: Professional Member

Skills

Research: Safe RL, weak supervision / sparse feedback, planning, constraint learning, optimization

Programming: Python, PyTorch, NumPy/Pandas, Matplotlib, SQL

Tools: Linux, Git, \LaTeX

Languages: English (fluent), Mandarin Chinese (fluent)

Competitive Programming: Google Hash Code 2017 — Rank #46 (UK), #627 (EMEA)