Yiming Xiong

yimingx920@gmail.com | yimingxiong.me
Leah Xiong | zksha | yimingxiong
Vancouver, British Columbia, Canada

Education

- **University of British Columbia** *Sep. 2025 - Jun. 2026*
  Master of Data Science (Computational Linguistics) Vancouver, Canada
  - GPA: 4.33/4.33
- **Chinese University of Hong Kong, Shenzhen** *Sep. 2020 - Jun. 2024*
  Bechelor of Data Science and Big Data Technology Shenzhen, China
  - GPA: 3.8/4.0

Research Interest

End-to-end autonomous scientific research, with a focus on model self-improvement through reflection, continual learning, and experience-driven training, and on bringing LLMs and AI scientists into the physical world.

Research Experience

- **RealLabBench: Towards Autonomous Agentic Robot for Science** *Present*
  Collaboration with Tsinghua University | Ongoing research
  - Developing Anything-to-Sim, an end-to-end pipeline for creating simulation-ready laboratory equipment assets from technical drawings of specialized scientific instruments.
  - Composing these assets into realistic, interactive biology and chemistry laboratory scenes, with equipment geometry and physical states for simulating experimental workflows.
  - Designing RealLabBench to evaluate language models and robots on end-to-end, long-horizon scientific workflows in realistic simulated laboratories, integrating reasoning, instrument operation, and physical interaction.
- **ALMA: Learning to Continually Learn via Meta-learning Agentic Memory Designs [🌐]** *Sep. 2025 – Feb. 2026*
  University of British Columbia | Collab.: Shengran Hu, Jeff Clune
  - **Accepted at NeurIPS 2026**; **Outstanding Paper**, ICLR 2026 Recursive Self-Improvement Workshop (Top 3 out of 110); **Best Paper**, ICLR 2026 MemAgent Workshop (Top 4 out of 106).
  - Formulated a structured search space for external memory designs by analyzing existing agentic memory mechanisms, enabling meta-agents to discover flexible memory architectures.
  - Designed a meta-agent discovery pipeline that enables agents to ideate, implement, and evaluate memory designs under static and dynamic memory settings.
  - Demonstrated that meta-agents can discover complex memory designs beyond human-designed SOTA baselines, contributing evidence toward AI-driven automated research and the evolution of agent intelligence.

Work Experience

- **Kimi / Moonshot AI [🌐]** *Jun. 2026 – Present*
  Post-training Research Intern, Foundation Model Team Beijing, China
  - Analyzing foundation model behavior and failure modes to guide post-training data design and reinforcement learning pipeline development for self-evolving capabilities.
  - Responsible for designing RL pipelines and constructing post-training data to support iterative self-evolution and automated research capabilities in foundation models.
- **DataStory [🌐]** *Sep. 2024 - Aug. 2025*
  Machine Learning Engineer Guangzhou, China
  - Online evolution of tool usage and management strategies for social media deep research agents. Designed and implemented an online-evolving tool archive by allowing agents to alternate tools parallel while interact with users. Enabled agents to dynamically adapt tool selection and planning strategies based on historical interactions, improving multi-hop reasoning robustness (+15%) and reducing redundant tool calls.
  - Led large-scale reward model fine-tuning for LLM domain-specific reasoning improvement, designing hybrid reward mechanisms combining text-based and verifier signals using historical agent interaction data. Enabled domain-specific research and reflection strategies for planning and self-evaluation agents, improving reasoning performance by 11% in quality of generated report.
  - Provided technical consulting and research support for cross-functional business teams, assisting in AI system integration, and domain-specific deployment optimization across multiple product scenarios.

PATENTS AND PUBLICATIONS C=CONFERENCE, J=JOURNAL, P=PATENT, S=IN SUBMISSION, T=THESIS, R=TECHNICAL REPORT

[C.2] **Xiong, Y.**, Hu, S., Clune, J. (2026). **Learning to Continually Learn via Meta-learning Agentic Memory Designs**. In *Advances in Neural Information Processing Systems (NeurIPS 2026)* (accepted).

[R.1] Kimi Team et al. (including **Yiming Xiong**) (2026). **Kimi K3: Open Frontier Intelligence**. *arXiv preprint arXiv:2607.24653*.

[C.1] Yuan, Y., Wang, W., Guo, Q., **Xiong, Y.**, Shen, C., He, P. (2024). **Does ChatGPT Know That It Does Not Know? Evaluating the Black-Box Calibration of ChatGPT**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. ELRA and ICCL. May 2024, Torino, Italy. DOI: 10.18653/v1/2024.lrec-main.462

SKILLS

- **Programming Languages:** Python, Node.js, SQL, Bash, Stata, R, Matlab
- **Frameworks & Libraries** PyTorch, Tensorflow, Transformers, OpenRLHF, VeRL, LLaMA-Factory, VLLM, LangGraph, FastAPI, Ray, SSE
- **DevOps & Tools:** Linux, Docker, Git
- **Research Skills:** Meta-learning for agentic system, Agentic system evolution, reinforcement learning for reasoning improvement, large language model reliability and calibration analysis, multi-agent reasoning systems, and scalable AI experiment design.