

SICONG JIANG

+1 404-277-4365 | [Homepage](#) | [Email](#) | [Github/JohnsonJiang](#) | [LinkedIn/Sicong](#) | [Google Scholar](#)

SUMMARY

Final-year McGill Ph.D. and Google DeepMind Research Intern specializing in **Recursive Self-Improvement (RSI)**, **Post-Training**, and **Automated Reward Modeling**. Recognized for pioneering **human-aligned reward models** and architecting **production-grade agentic data/evaluation pipelines** from scratch. First-author contributor to top AI venues (ICLR, EMNLP, ICRA, AISTATS) with **900+ citations**.

EDUCATION

McGill University

Ph.D in Center for Intelligent Machines

Jan 2022 – Present

Quebec, Canada

Georgia Institute of Technology

Master of Electrical and Computer Engineering

Sep 2019 – Dec 2021

Georgia, USA

Northeastern University

Bachelor of Electrical Engineering

Sep 2015 – June 2019

Shenyang, China

EXPERIENCE

Google DeepMind

Research Intern, Working on RSI

Mar 2026 – Present

London, UK / Toronto, Canada

- **Self-Evolving Agents:** Engineered an execution-grounded RL architecture for **autonomous self-improvement**, enabling Gemini Flash-tier models to **outperform Pro-tier baselines** on competitive coding benchmarks.
- **Test-Time Compute Distillation:** Designed an execution-consistency verification pipeline, internalizing **test-time scaling compute** into permanent policy weight updates through continuous RL.
- **Harness-Data Co-Evolution:** Architected an automated harness pipeline to actively harvest **high-hardness RL data** and edge cases, driving continuous model gains via the co-evolution of evaluation harnesses and trajectory data.
- **RL Diagnostics & Reward Hacking:** Developed an automated pipeline to identify and mitigate **reward hacking loops**; enhanced training stability and policy performance by filtering anomalous agentic rollouts in the RL loop.

Abaka AI

Director of Research, Founding Scientist

Aug 2025 – Feb 2026

Palo Alto, California

- **Research (Benchmarking & Evaluation):** As a founding member of the Research team, led benchmarking and evaluation for agentic and multimodal LLMs; led **EditReward** (ICLR'26) and co-developed large-scale benchmarks including **SuperGPQA** (NeurIPS'25), **ChartNet** (CVPR'26), **EgoTL** (CVPR'26), and **VeriWeb**.
- **Advanced Dataset & Pipeline Design:** Led multiple **zero-to-one pipeline builds**—architecting and deploying high-difficulty datasets and production pipelines from scratch across coding, IMO-level math, multimodal data, agentic trajectories, and **RL environments**.
- **Frontier AI Deployment:** Delivered mission-critical datasets and pipelines deployed directly for foundation model **post-training/evaluation** across multiple **frontier AI labs**.

Mercor

Applied Scientist Intern

May 2025 – Aug 2025

Remote

- **Data Quality & Evaluation:** Designed **multi-stage QA** for large-scale LLM training annotations, improving consistency via calibration, automated checks, and targeted reviewer feedback loops.
- **Metrics & Analytics:** Built labeler-performance metrics and monitoring to detect systematic errors and guide guideline iterations for higher-quality training data.

SELECTED PUBLICATIONS

10+ top-tier conference papers with 900+ citations. * indicates equal contribution. For full list, visit [Google Scholar](#).

- **[ICLR 2026]** Wu K*, Jiang S*, Ku M, et al. *EditReward: A Human-Aligned Reward Model for Instruction-Guided Image Editing. (Human-Aligned Reward Modeling)*

- [EMNLP 2025] Qian K*, **Jiang S***, Zhong Y, et al. *AgentThink: A Unified Framework for Tool-Augmented Chain-of-Thought Reasoning in Vision-Language Models for Autonomous Driving*. (Agentic Post-Training & Tool Use)
- [ICRA 2026] Luo Z*, **Jiang S***, Qian K*, Huang Z, Miao J, et al. *FASIONAD+: Enhanced Safety in Autonomous Driving with Adaptive Feedback*. (Autonomous Systems & VLA)
- [AISTATS 2025] Liu F*, **Jiang S***, et al. *Adversarial Vulnerabilities in Large Language Models for Time Series Forecasting*. (LLM Robustness & Red-Teaming)
- [ICCV-W 2025] **Jiang S***, Huang Z*, Qian K*, et al. *A Survey on Vision–Language–Action Models for Autonomous Driving*. (Vision–Language–Action Ecosystem)

SKILLS

Languages & Core: Python, C/C++ , SQL, Shell/Bash, Git, Linux

Frameworks & Systems: PyTorch, veRL, Ray, vLLM, SGLang, DeepSpeed, HuggingFace (Transformers/TRL)

Post-Training & Alignment: Multi-Step RL, Scalable RL (PPO, DPO, GRPO), PRM/ORM, SFT, LoRA, Self-Distillation

Data Strategy & Agents: Synthetic Data Pipelines, Agentic Trajectory Synthesis, Multimodal Evaluation, Planning

SELECTED PROJECTS & OPEN-SOURCE IMPACT

Awesome-VLA4AD: Vision–Language–Action Ecosystem GitHub ★ 613
Creator & Primary Maintainer, Featured by Leading Tech Media

- Architected and maintain the premier open-source taxonomy tracking 200+ foundation models, benchmarks, and simulators for Vision-Language-Action (VLA) in embodied systems.
- Widely recognized across industry and academia as the standard community reference; covered in-depth by major AI media.

Open-Source Multimodal Datasets & Evaluators Hugging Face 20K+ Downloads
Core Architect & Benchmark Lead

- Released high-impact multimodal datasets on Hugging Face: **ChartNet** (14K+ downloads, MIT News Featured), **EditReward** (2K+ downloads), and **VeriWeb** (5K+ downloads).
- Engineered reusable evaluation harnesses adopted by external research teams at **Meta**, **Google DeepMind**, **Adobe**, and **Salesforce**.

Foundation Model Red-Teaming & Robustness Suite
Project Lead, Collaborated with Amazon Chronos & Nixtla

- Developed open-source gradient-free adversarial evaluation suites for foundation time-series forecasters, raising MAE/MSE by up to **45%** under constrained **2%** perturbations.
- Collaborated directly with foundation model teams (**Amazon Chronos**, **Nixtla TimeGPT**) to establish robustness baselines.

AWARDS

2021 – 2024 McGill Engineering Doctoral Awards (MEDA), McGill University

2021 TISED Doctoral Recruitment Awards (DRA), McGill University

2019 Outstanding Graduates of Liaoning Province

2019 Most Influential Graduates, Northeastern University

2018/2019 Northeastern University 1st Class Academic Scholarship

2017 National 1st Prize Award, China Undergraduate Mathematical Contest in Modeling

2016 Northeastern University 2nd Class Academic Scholarship

ACADEMIC LEADERSHIP & SERVICE

Workshop Organizer: CVPR 2026 (Video Generative Models: Benchmarks & Evaluation); ICCV 2025 (Memory and Vision); COLM 2025 (AI Agents: Capabilities and Safety).

Conference Reviewer: NeurIPS, ICLR, AISTATS, CVPR, ICCV, COLM, EMNLP, AAAI, IROS, ICRA.