

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 focusing on **AI Agents, Multimodal Model Evaluation, and Reward Modeling**. Recognized for developing **human-aligned reward models** and architecting **large-scale data/evaluation pipelines** from scratch. First-author contributor to top AI venues (**ICLR, EMNLP, ICRA, AISTATS**) with deep engineering proficiency in modern **LLM training paradigms**. **Actively seeking full-time industry roles** starting in **Fall 2026**.

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, Autonomous Agent Team

Mar 2026 – Sep 2026

London, UK

- **Self-evolving agents:** Developed a **Multi-Agent Grounded Debate Pipeline** that enables agents to **improve without verified ground truth** by pairing code generators with counter-case executors, helping **Gemini Flash** models consistently outperform **Gemini Pro** baselines on competitive programming benchmarks.
- **Execution-grounded rewards:** Built a **Dual-Reward System** combining semantic checks with sandbox execution, reducing hallucination and local-optima loops through compilation-grounded feedback.
- **Agentic post-training:** Developed a **multi-step RL** with online self-distillation pipeline to internalize execution-verified debate trajectories and improve zero-shot self-correction.

Abaka AI

Director of Research, Founding Member of Technical Team

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**.
- Delivered **datasets and pipelines** used directly for **model training/evaluation** for 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 700+ 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. (Multi-modal 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-train)*

- [ICRA 2026] Luo Z*, **Jiang S***, Qian K*, Huang Z, Miao J, et al. *FASIONAD+: Enhanced Safety in Autonomous Driving with Adaptive Feedback*. (VLM&VLA)
- [AISTATS 2025] Liu F*, **Jiang S***, et al. *Adversarial Vulnerabilities in Large Language Models for Time Series Forecasting*. (LLM Safety)
- [ICCV-W 2025] **Jiang S***, Huang Z*, Qian K*, et al. *A Survey on Vision–Language–Action Models for Autonomous Driving*. (VLM&VLA)

SKILLS

- **Languages:** Python, SQL, Java, MATLAB
- **Frameworks & Systems:** PyTorch, HuggingFace, veRL, TensorFlow, Scikit-learn, XGBoost
- **Post-Training & Alignment:** Supervised Fine-Tuning (SFT), Parameter-Efficient Tuning (LoRA), Scalable RL (PPO, DPO, GRPO), Reward Modeling
- **Data Strategy & Agents:** Synthetic Data Pipelines, Agentic Trajectory Synthesis, Multimodal/Agent Evaluation, Long-Horizon Planning

SELECTED PROJECTS

EditReward: A Human-Aligned Reward Model for Image Editing

Project Lead, Collaborated with University of Waterloo

- Spearheaded a **human-aligned reward model** for instruction-guided image editing; engineered scalable pipelines to construct **EditReward-DATA** (200K+ preference pairs) and **EditReward-Bench**.
- Achieved broad industry adoption, with the benchmark and dataset utilized by leading research teams at **Meta, Google DeepMind, Apple, Adobe, Salesforce**; boosted GenAI-Bench by +23% over the Qwen2.5-VL baseline via reward-driven data curation.

Agent-Think: Tool-Augmented Reasoning in VLMs for Autonomous Driving

Project Lead, Collaborated with Xiaomi Inc. & Tsinghua University

- Proposed **Agent-Think**, a VLM framework for autonomous driving that combines structured CoT with **tool-augmented** reasoning; built a pipeline to auto-generate self-verified VQA with CoT & tool-use.
- Fine-tuned Qwen2.5-VL-7B with **SFT + GRPO** (LoRA), improving answer accuracy up to +54% and reasoning score up to +33% on public benchmarks.

Adversarial Attacks in LLM-based Time-Series Forecasting

Project Lead, McGill University

- Designed a **gradient-free, black-box attack** suite for LLM-based forecasters, raising MAE/MSE up to 45% with perturbations limited to 2% of the input-signal mean.
- Curated an automated evaluation pipeline and collaborated with several time-series foundation model teams (**Amazon Chronos, Nixtla TimeGPT**).

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

OTHER EXPERIENCE

Open-Source Contributor: **ChartNet** (65K total downloads on Hugging Face; **MIT News Featured**), **EditReward** (4K monthly downloads on Hugging Face), and **VeriWeb** (2K monthly downloads on Hugging Face).

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.