



Jiadong Guo

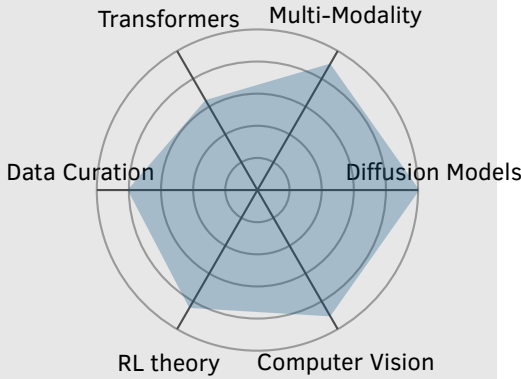
LLM Research Engineer

- July 1992
- Bay Area, CA, US
- <http://jd-unleashed.com>
- jiadong.guo@outlook.com

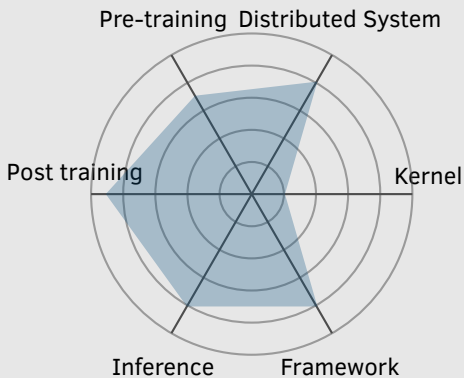
About me

Background in Multimodal Understanding and Generation, now focused on solving engineering problems at scale with RL.

Knowledge



Engineering



(*)[The scale is from 1 (Fundamental Awareness) to 5 (Domain Expert).]

Education

- 2011-2015 B.Sc. in Mechanical Engineering RWTH Aachen, Germany
- 2015-2018 M.Sc. in Robotics, Systems and Control Specializing in vision, learning and control. ETH Zurich, Switzerland

Experience

- Oct 2025 Member of Technical Staff Microsoft AI Infra, US
 - Scaling, Fault Tolerance, Efficiency, Observability in RL
 - Support all MAI models' post-training
 - Lead VLM/Omni RL training infra
- Jan 2025 Research Engineer Meta Superintelligence Lab, US
 - Post Training for Llama 4 multi-image reasoning capability
 - Post Training captioner for t2i/t2v generation (Emu/MovieGen)
 - Synthetic image composition data for editing model (EmuEdit)
 - Distributed RL training infrastructure for agentic multimodal tool use
- Dec 2023 Computer Vision Engineer Meta Monetization GenAI, US
 - Pre-train and DPO finetune a text-to-image model
 - Develop visual text rendering capability with 119% adoption gain
 - Serving the model to meta-wide advertisers
 - Maintaining internal media generation training framework
- Jun 2020 Computer Vision Engineer Meta Oculus, Switzerland
 - Synthetic data for eye tracking model training using 3D rendering engines, domain adaptation, GAN and neural radiance fields.
- Apr 2019 Robot Perception Engineer Sevensense Robotics, Switzerland
 - Perception and navigation modules of autonomous ground industrial vehicles with LiDAR, depth sensor and stereo cameras

Open Source

- Aug 2025 Sglang Developer Network RL Systems
 - 7x faster weight transfer for 1T OSS model on RL framework
 - Several bug fixes
 - Upstream tau-bench agentic RL example on slime

Publications

- 2019 Local descriptor for robust place recognition using LiDAR intensities ICRA/RA-L 2019, 1st author
- 2019 Robust photogeometric localization over time for map-centric loop closure RA-L 2019, 4th author
- 2022 Compilergym: Robust, performant compiler optimization environments for ai research CGO 2022, 3rd author
- 2023 Learning compiler pass orders using coreset and normalized value prediction ICML 2023, 6th author
- 2026 Scaling Sequence-to-Sequence Generative Neural Rendering ICLR, 4th author
- 2026 Mixture of States: Routing Token-Level Dynamics for Multimodal Generation CVPR, 10th author