Rodney Lafuente-Mercado

New York, NY | (605) 202-0853 | rodneyslafuente@gmail.com | rlafuente.com | Active TS Clearance

Experience

Scale AI Machine Learning Engineer, Public Sector — New York, NY | Feb 2026 to present

- Own agent evaluation for Donovan, Scale's national-security LLM platform, measuring whether agents can complete real government workflows for public-sector customers.
- Led a public benchmark of LLMs on military planning, built with subject-matter experts: 20 frontier models tested on 1,257 expert-written requirements, with the best model meeting 62% of them.
- Automated the grading with an LLM judge that scores as consistently as credentialed military planners.
- Researched RL fine-tuning of LLMs with GRPO and LoRA: found why token-level reward weighting breaks down and built a simpler fix (sole-author paper, ICML 2026 workshop).

New York University Research Collaborator with Pavel Izmailov — 2026 to present

- Built a pipeline that uses frontier LLMs and the Lean proof assistant to generate training data with no human annotators: 10,634 proofs and 152K labeled steps.
- Trained proof-checking models from 0.6B to 8B parameters with GRPO on H100s, using VERL and vLLM.
- Raised step accuracy on held-out proofs from 48% to 75%, beating the 60% from OpenAI's human-labeled PRM800K dataset (first author, NeurIPS 2026 workshop).

MIT Lincoln Laboratory Assistant, then Associate Staff — Cambridge, MA | Jun 2023 to Jan 2026

- Wrote autonomy software that flew on the X-62A VISTA, a modified F-16, in DARPA flight tests at Edwards AFB, on the program that flew the Secretary of the Air Force.
- Took a new MIT CSAIL sequence model from research paper into DARPA flight simulation, its first use in online RL, and trained policies from toy tasks up to simulated aerial combat.
- Built RL infrastructure for 100+ node clusters, including an open-source distributed RL library (sole-author paper, IEEE HPEC 2025) and containerized Ray on Slurm (IEEE HPEC 2024).

Amazon Web Services Software Engineer Intern — Seattle, WA | Jun 2021 to Aug 2021

- Shipped production API changes to AWS Signer, Amazon's code-signing service.

Selected Publications

- LeanPRM: Lean-Grounded Process Supervision for Natural Language Proof Verification. First author. NeurIPS 2026 MATH-AI Workshop.
- MA-Bench: Measuring Operational Planning in Language Models. First author. Under review.
- ARCA: Adapter-Residual Credit Assignment When Token Signals Degenerate. Sole author. ICML 2026 DEMO Workshop.
- Adaptive Policy Synchronization for Scalable Reinforcement Learning. Sole author. IEEE HPEC 2025.

Education

The University of Texas at Austin M.S. in Computer Science — Expected May 2027

Thesis on tracing language-model behavior back to RL training data, advised by Prof. Leqi Liu.

Harvard University B.A. in Applied Mathematics — December 2022

Skills

Tools: Python, C++, CUDA C++, PyTorch, VERL, vLLM, Ray, Slurm, Docker, AWS, Lean 4
Methods: reinforcement learning (PPO, GRPO), LLM fine-tuning, LLM evaluation, distributed training
Languages: English, Spanish (fluent)