

Rodney Lafuente-Mercado

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

Research interests: Reinforcement learning, LLM reasoning and verification, credit assignment, and robust autonomous systems.

EDUCATION

The University of Texas at Austin

M.S. in Computer Science, part-time

Current research with Leqi Liu on reinforcement learning and credit assignment.

Austin, TX

Expected May 2027

Harvard University

B.A. in Applied Mathematics

Cambridge, MA

December 2022

RESEARCH AND PROFESSIONAL EXPERIENCE

Scale AI

Machine Learning Engineer, Public Sector

New York, NY

Feb 2026 - Present

- Create evaluations for large language models using input from subject-matter experts (SMEs), benchmarking emerging model capabilities on domain-specific tasks relevant to national security.
- Guide subject-matter experts in developing MA-Bench, a benchmark measuring operational planning in language models; coauthor the paper on this benchmark developed entirely at Scale AI.
- Developed ARCA, using adapter hidden-state residuals to assign token-level credit during parameter-efficient LLM reinforcement learning without a learned reward model or value head.
- Studied degeneracy in token-level credit signals under LoRA through controlled GRPO experiments with matched rollout budgets; sole-author paper accepted to the ICML 2026 DEMO Workshop.

LLM Reasoning and Verification

Research collaboration with Pavel Izmailov and Salman Rahman, NYU

Current collaboration

- Build pipelines for synthetic data generation, RL training, and evaluation of natural-language proof verifiers using step-level correctness labels derived from formal proof attempts checked in Lean.
- Distinguish mathematical reasoning errors from formalization failures, reviewing correspondence between informal proof steps and formal traces before constructing step-level training labels.
- Train verifiers that assess informal mathematical proofs without a proof assistant at inference; evaluate accuracy on held-out theorems and transfer across external proof-verification benchmarks.

MIT Lincoln Laboratory

Associate Staff

Assistant Staff

Cambridge, MA

Aug 2025 - Jan 2026

Jun 2023 - Jul 2025

- Developed and evaluated reinforcement learning systems for aircraft autonomy under DARPA Air Combat Evolution (ACE), including contributions flown on the X-62A VISTA research aircraft.
- Worked within Daniela Rus's CSAIL group while employed by Lincoln Laboratory; developed recurrent and state-space models from research prototypes through simulated aerial combat.
- Coauthored *Temporal Range* on policy sensitivity to past observations and *When Sensors Fail* on robust PPO, studying memory use and robustness under partial observability and sensor drift.
- Built RL infrastructure for 100+ node Slurm clusters; developed ClusterEnv and adaptive policy synchronization, and co-developed Syndeo to run portable, containerized Ray workloads on Slurm.

PUBLICATIONS AND PREPRINTS

MA-Bench: Measuring Operational Planning in Language Models

Rodney Lafuente-Mercado, Robert Sincero, Phil Sage, and Yuan (Emily) Xue.

Paper in preparation. Research conducted at Scale AI.

LeanPRM: Lean-Grounded Supervision for Natural Language Proof Verification

R. Lafuente-Mercado, S. Rahman, J. Shen, J. Shiffer, S. Gabriel, and P. Izmailov.

NeurIPS 2026 MATH-AI Workshop, accepted poster.

OpenReview

ARCA: Adapter-Residual Credit Assignment When Token Signals Degenerate

R. Lafuente-Mercado.

ICML 2026 DEMO Workshop.

arXiv:2606.00257

When Sensors Fail: Temporal Sequence Models for Robust PPO under Sensor Drift

K. Vogt-Lowell, T. Tsiligkaridis, **R. Lafuente-Mercado**, S. Ghatti, S. Gao, M. Zitnik, and D. Rus.

ICLR 2026 CAO Workshop.

arXiv:2603.04648

Temporal Range: Measuring History Use in Sequence Models

R. Lafuente-Mercado, D. Rus, and T. K. Rusch.

Preprint, 2025.

arXiv:2512.06204

Adaptive Policy Synchronization for Scalable Reinforcement Learning

R. Lafuente-Mercado.

IEEE High Performance Extreme Computing (HPEC), 2025.

arXiv:2507.10990

Syndeo: Portable Ray Clusters with Secure Containerization

W. Li, **R. Lafuente-Mercado**, J. D. Peña, and R. E. Allen.

IEEE High Performance Extreme Computing (HPEC), 2024.

arXiv:2409.17070

EARLIER EXPERIENCE

Harvard Medical School

Undergraduate Researcher

Boston, MA

Dec 2021 - May 2022

- Developed a pipeline combining machine learning and image processing to analyze SARS-CoV-2 nucleocapsid images; trained laboratory staff to use the workflow without requiring ML expertise.

Amazon Web Services

Software Engineer Intern

Seattle, WA

Jun 2021 - Aug 2021

- Implemented AWS Signer API changes supporting fine-grained signing configurations; integrated the changes into production using the AWS Cloud Development Kit (CDK) to manage infrastructure.

TECHNICAL SKILLS

Programming: Python, C++, CUDA C++.

Research and infrastructure: PyTorch, Ray, Slurm, Lean 4, AWS, Git.

Languages: English; Spanish (fluent).