

Romir Patel

United States - Citizen | P: +1 510-992-0406 | rompr18@gmail.com

WORK EXPERIENCE



Bespoke Labs

Mountain View, CA

Member of Technical Staff: Machine Learning Research Engineer - Post-Training

July 2026-Present

Machine Learning Research Engineer - Intern

March 2026-July 2026

- Engineering & led research on synthetic world creation for RL environments, generating consistently low-scoring tasks for SOTA models (Claude Mythos/Fable 5), reducing cost by 10x (from 100,000 to 10,000 per world) for complex RL environment and task world creation
- Led & built end-to-end RL environment generation infrastructure for coding agents for multiple frontier labs producing 5,000+ OTS training and evaluation environments across Python, Go, JavaScript/TypeScript, and Java, responsible for all aspects spanning environment construction, reward design, task generation, and automated validation infrastructure at production scale.
- Designed automated multi-agent episode synthesis generating realistic developer work artifacts (branches, PRs, review threads, failed CI runs, competing fixes) to create high-fidelity RL training signal grounded in real codebases
- Engineered a scalable task hardening pipeline that systematically removes answer leakage, injects adversarial candidate solutions, and enforces cross-source evidence synthesis, producing environments difficult enough to induce consistent failures in frontier-class coding agents.
- Developed end-to-end task validation with Docker-based hermetic builds, pre-patch failing tests, and post-patch oracle verification, enforcing strict correctness guarantees across 3,000+ generated environments with zero tolerance for grader false positives.
- Diagnosed and resolved grader brittleness where semantically correct agent rollouts failed due to surface-level matching errors; redesigned grading logic to evaluate solution correctness at the semantic level while preserving hidden-test rigor.



AgentHub

San Francisco, CA

Machine Learning Research Engineer - Intern

January 2026-March 2026

- Architected end-to-end LLM evaluation infrastructure with SQL-based ground truth validation, processing 1M+ rows across production domains; reduced evaluation set creation from days to minutes while maintaining 100% validation accuracy
- Built closed-loop synthetic data refinement pipeline — LLM generates, SQL validates against ground truth, system provides corrective feedback, LLM regenerates — with statistical validation (KS/Chi-square) enforcing configurable 75–95% distribution alignment
- Designed multi-agent orchestration framework coordinating planner, executor, and validator sub-agents via DAG-based task decomposition with a 6-tier complexity curriculum from single-table lookups to multi-join relational reasoning
- Developed async database connector framework supporting PostgreSQL, MySQL, and SQLite with provider auto-detection across 8 cloud platforms; connection pooling and structured error attribution reduced user debugging time ~80%
- Implemented multi-modal agent reasoning combining vision encoders with SQL execution for document-grounded database queries, reducing context window requirements 40% through structured image-to-schema mapping



Massachusetts Institute of Technology (MIT), Computer Science & Artificial Intelligence Laboratory

Cambridge, MA

Machine Learning Researcher - Post Training

July 2025-Present

- Architected multimodal autoformalization agents that process LaTeX and diagram inputs jointly, achieving 68% type-checking success rate through vision-language model integration with incremental verification
- Reduced LLM proof hallucination 30%+ and invalid proof generation 35% by developing bidirectional cycle consistency validation across 1000+ Lean theorems and miniF2F benchmark
- Increased formal proof generation throughput 5x via incremental type-checking on 8B-parameter LLMs; built comprehensive benchmarking suite evaluating 8 code generation models across 1000+ Lean translation tasks
- Published: [Improving Autoformalization via Cycle Consistency and Incremental Type-Checking Using Language-Model Probabilistic Programs](#) — NeurIPS 2025 (Math-AI)



California Institute of Technology (Caltech)

Pasadena, CA

Machine Learning Researcher - Post Training

August 2025-March 2026

- Training and fine-tuning frontier-scale generative models (LLaDA-8B-Instruct, LLaDA-MoE-7B) using reinforcement learning under Professor Anima Anandkumar, ex-Sr. Director of AI @ NVIDIA
- Optimized distributed RLHF training across 32 GPUs using DeepSpeed ZeRO-3 and gradient checkpointing, reducing alignment compute cost 2.5x and memory footprint 4x while preserving training throughput

Built agent fine-tuning pipeline combining supervised learning on 50K+ theorem proofs with RL from automated verifier feedback, improving multi-step mathematical reasoning accuracy 35% across 5K test cases



Massachusetts Institute of Technology(MIT) | McGovern Institute for Brain Research

Cambridge, MA

Machine Learning Researcher - Applied & Post Training

July 2025-June 2026

- Led large-scale ML study (N=23,760) developing novel ensemble stacking methods that achieved $R^2=0.50-0.60$ for predicting depression and anxiety across developmental trajectories, combining transformer-based models (TabPFN), gradient boosting (CATBoost, XGBoost), and Random Forests with rigorous nested cross-validation
- Engineered feature analysis pipeline across 1,078+ biopsychosocial variables spanning 30 theoretical domains, applying PCA to 420+ neuroimaging features and validating findings via ML, Bayesian Mixed Models, and Cross-Lagged Panel Networks
- Built model evaluation framework incorporating SHAP values, permutation importance, Bayesian uncertainty quantification, and bootstrap stability analysis (2,500 iterations) following consensus ML best practices for scientific reproducibility
- Paper conditionally accepted to *Nature*



Massachusetts Institute of Technology(MIT) | School of Engineering

Cambridge, MA

Machine Learning Researcher - Applied

November 2024-January 2026

- Developed hybrid neural-symbolic agent for autonomous PDE solver selection and hyperparameter optimization, achieving 25–100x speedup over baselines through learned control policies across 3 benchmark systems
- Implemented custom CUDA kernels for PDE solver acceleration, achieving 15x speedup over NumPy through optimized memory coalescing and shared memory utilization
- Engineered novel neural architecture achieving $<6\%$ L2 error on Navier-Stokes uncertainty quantification; reduced stochastic MPC computational cost 33x on 2D systems via polynomial chaos expansion
- **Paper: *Physics Informed Neural Networks with Model Predictive Control* — submitted to Automatica**



TextArena.ai

Remote

Machine Learning Researcher - Post Training

November 2024-December 2025

- Co-developed open-source agentic evaluation framework enabling 100K+ autonomous agent interactions across 74+ strategic reasoning environments with self-play RL training pipelines and Gym-style interface
- Built TrueSkill-based evaluation system ranking 216+ models across theory-of-mind, negotiation, and deception tasks; 50,000+ environments played by users and models within 5 months of launch
- Designed automated benchmarking infrastructure with statistical significance testing (bootstrap CI, effect sizes) comparing agent architectures across **1000+ evaluation episodes**
- Raised \$350K+ in funding from OpenRouter, Plastic Labs, Anthropic, and Amazon; project recognized by Andrej Karpathy and acknowledged by Demis Hassabis and Elon Musk
- Published work & progress in a [IBM Research Blog](#)

EDUCATION

Carnegie Mellon University | School of Computer Science

Pittsburgh, PA

Bachelors of Science, Computer Science and Artificial Intelligence

Basha High School

Chandler, AZ

High School Diploma. AP Scholar with Distinction. National Merit Scholar.

Focus in CS, HOSA VP, Speech & Debate Captain, DevOps Co-President & International Competitor

ACTIVITIES

The KeelWorks Foundation

Oak Harbor, WA

Board Member

- Became the foundations **youngest ever board member** in its 20+ year history
- Currently working on the development of a startup accelerator branch, The Cube, an incubator for pre-seed tech startups

SKILLS

Agentic AI & LLM Stack: LangChain, AutoGPT, LangGraph, OpenAI Agents, Claude API, vLLM, TensorRT-LLM

Agent Frameworks: ReAct, ReWOO, Tree-of-Thoughts, Chain-of-Thought prompting, function calling

Multimodal: CLIP, LLaVA, vision-language model fine-tuning, OCR integration