Ohm Patel

San Francisco, CA | ohmpatel1212@gmail.com | ohm.cx | GitHub | LinkedIn

AI engineer focused on physical reinforcement-learning environments for scientific discovery. Background in LLM post-training, data quality, evaluation, and inference systems.

Research & scientific environments

Scientific environments - Built three synthetic agent environments: causal circuits, reaction kinetics, and adaptive spectroscopy. Agents choose measurements under a budget and predict unseen outcomes. Prototypes; not yet validated in real labs.

ProteoBench - Developing a benchmark for agents analyzing proteomics data, with deterministic graders and a paper-mining and expert-review pipeline for dataset construction.

RL environment atlas - Published a sourced map of 190 environment types, covering tasks, data, and rewards. Wrote A Mechanistic Fallibilist on agents learning through experiments.

Experience

Hooglee / Member of Technical Staff — Aug 2025 - Jul 2026

- Early engineer at Sebastian Thrun and Eric Schmidt's AI lab, working across LLM post-training, data pipelines, and video evaluation.
- Post-trained Qwen-family models with reinforcement learning; built environments and evaluation harnesses for video-editing quality.
- Owned annotation quality and built a vendor pipeline producing approximately 100,000 labeled training and evaluation examples.
- Designed video-understanding evaluations for clip ordering, speech cut-offs, and mechanical editing errors.

Microsoft AI / Software Engineering Intern — May 2025 - Aug 2025

- Built a simulator for disaggregated LLM inference, modeling prefill/decode and request batching to compare GPU architectures across workload traces.
- Measured time-to-first-token and time-per-output-token to guide GPU selection for different workloads.

VigilAI / Founder — Jan 2024 - Jul 2025

- Built LLM tools that draft police reports using body-camera footage. Backed by Fusen Capital; participated in Georgia Tech CREATE-X.

UCF / Lockheed Martin / Research Engineer — Sep 2024 - Dec 2024

- Built LoRA fine-tunes for LLM-driven simulations in a Lockheed Martin research engagement.

Earlier: CSAA Insurance Group, Software Engineer Intern (May - Aug 2024): AWS claims automation. Career Fair Plus, Software Engineer Intern (Mar 2023 - Feb 2024): React tools for online career fairs.

Selected projects

Amadi - Engineering contributor to custom social feeds refined through chat.

Spydr - Knowledge and memory tools; 1,000+ users. Therapute - Computer-vision feedback for physical therapy; Google Solution Challenge top ten.

Education & skills

B.S. Computer Science, University of Central Florida | May 2025

ML / systems: LLM post-training, RL environments, evaluation, data pipelines, disaggregated inference, GPU simulation. Tools: Python, PyTorch, C, TypeScript, React / Next.js, AWS.