

# LILY ZHANG

San Francisco Bay Area | [lilyzhng.ai@gmail.com](mailto:lilyzhng.ai@gmail.com) | [linkedin.com/in/lilyzhng](https://www.linkedin.com/in/lilyzhng) | [lilyzh.ng](https://lilyzh.ng) | [github.com/lilyzhng](https://github.com/lilyzhng) | [@lily\\_gpupoor](https://twitter.com/lily_gpupoor)

## STATEMENT

---

Applied research lead with 7 years building o-to-1 technical programs across post-training, AI training data, and RL environments. I have transformed research orgs from R&D to product-oriented applied ML engineering, leading global ML teams across 4 countries and shipping research into production platforms. I have 8 publications at CVPR, NeurIPS, and ICCV, focusing on post-training, synthetic data, and evaluation.

## HIGHLIGHTS

---

- **Built RL Environments and Eval Benchmarks:** Built SuperGeneral, compositional tool environments for long-horizon agents on APEX benchmark; Designed verifiable failure taxonomy for RL reward shaping (tested on Terminal Bench &  $\tau^2$ -bench). Built LosslessBench evaluating inference acceleration impact on model behaviors.
- **Large-Scale Training & Resource Management:** Owned compute allocation for 1000+ GPUs across competing teams, designed a dynamic priority scheduling system for cluster sharing. Built SofaGenius, an automated post-training platform, productionized for internal ML experiments.
- **o-to-1 Program Delivery:** Decomposed a fuzzy research goal into a phased roadmap with milestones, owners, and review checkpoints. Stood up a scout team across perception, mapping, ops, and labeling. AI training data research shipped in Ford F-150 Lightning and Super Duty; auto-labeling pipeline delivered \$1M+ savings.
- **Program Leadership & Org Building:** Ran a global team of 10+ across 4 countries (US, India, Germany, Israel). Turned ambiguous customer data contracts into scoped deliverables with named owners, milestones, and a recurring PI-planning that aligned perception, ops, and labeling. At Latitude AI, drove the org-wide shift from traditional ML to agentic workflows. Skills/Tooling I built is now used by 76+ engineers.
- **Alignment Research & Governance Tooling:** [C-Guard: A Constitution-Grid Instrument for Data-Efficient RL Alignment](#) (accepted, COLM 2026 Efficient Reasoning Workshop). Turns the helpfulness/harmlessness tradeoff into measured, gated moves on a constitution grid. As research advisor at Alibaba Cloud, led [Eureka](#) from research to launching, monitored operational health, +16% demand fulfillment, -33% compute migration, 91% ops adoption. Poster @ Neurips 25 VLM4RWD, Oral @ DASFAA 2026.

## LEADERSHIP & SERVICE

---

- **NVIDIA GTC Conference.** Keynote on AI training data generation @ GTC 23. Panelist at AI/ML infra @ GTC 26
- **Industry Advisory Board:** Board Advisor for Texas A&M University, University of Delaware.
- **Research Advisor, AliCloud (2025)** Led the research direction for Eureka, AI infra agent post-training.
- **Program Committee:** NeurIPS, CVPR, ICRA, IROS, ICML, ICLR, ECCV, KDD, ICCV

## EXPERIENCE

---

### Technical Lead, AI & ML | Latitude AI

Feb 2024 - Present

- Managed compute allocation for 1000+ GPUs over 2 months to train 50000 models. Designed dynamic priority system with Dagster for cluster sharing across teams.
- Championed VLM/LLM adoption across the org. Built Alfa Agent for data selection; adopted org-wide with 76+ users.
- Formed a cross-functional scout team spanning perception, mapping, operations, and labeling; defined and phased the applied research roadmap for ML-based satellite-imagery segmentation, setting milestones for each phase.
- Designed agentic QA platform (Aqua), focusing on bring your own tool (BYO) for cost-effective auto QA and rework.

## Technical Lead, AI & ML | Ford Greenfield Labs

Jun 2022 - Jan 2024

- Influenced global team without authority, working with data-generation team of 10+ engineers across US, India, Germany, and Israel. Drove cross-org alignment and PI planning, turning customer data contracts into regular deliverables across perception workstreams.
- Transformed the research-scientist team into an applied ML engineering group: established production coding standards and mentored researchers into ML engineers with a product mindset and due diligence.
- Proposed and built auto-labeling pipeline replacing expensive manual annotation. Delivered 15K 3D annotations and \$1M+ in savings. Pipeline covers both perception labels and VLM-generated scene captions.
- Partnered with external groups, including Voxel51, Scale AI, Google, and NVIDIA on data and ML tooling.

## Sr. Research Scientist, AI & ML | Ford Greenfield Labs

Jun 2019 - May 2022

- Led full ML lifecycle from research prototype to production: synthetic data augmentation shipped in F-150 Lightning and Super Duty (Pro Trailer Hitch Assist). Bridged research and product teams.
- Pioneered AI-based 3D scene relighting, enabling reliable perception with cost-effective sensor configurations.

## PRODUCTS & OPEN SOURCE

---

- [SuperGeneral \(2026\)](#) Compositional tool environments for long-horizon agents on APEX benchmark. Tested 4 frontier models on tool use, composition, and creation.
- [RLVR Environment Design from Failure Taxonomy \(2026\)](#) What Can't Be Measured Can't Be Solved. Designed a verifiable failure taxonomy for RL reward shaping, benchmarked on Terminal Bench &  $\tau^2$ -bench.
- [GeniusTeam \(2026\)](#) Multi-agent orchestration. 4 specialized agents coordinate autonomously with file-based handoffs and worktree isolation.
- [SofaGenius Post-Train \(2026\)](#) Automated post-training platform. GRPO on 80B MoE models, cost-visible GPU launches on Modal, 7 anomaly detectors for model convergence.
- [Frontend Slides \(2026\)](#) Claude Code skill for HTML presentations. Reached tens of thousands of users, 25K+ github stars.

## SELECTED PUBLICATIONS

---

- [Eureka](#) (Oral DASFAA 2026, NeurIPS-W 2025) Agent that writes forecasting features as code.
- [LiHi-GS](#) (RAL/CVPR-W 2025) LiDAR-supervised Gaussian splatting for highway driving scene reconstruction.
- [SIMBAR](#) (CVPR 2022) Single image relighting for data augmentation. Diversifying ML training data.
- 7 granted US/CN/DE patents on synthetic data generation, augmentation, and curation.

## SKILLS

---

- **Technical Program Management:** scoping ambiguous problems, roadmap and milestone planning, cross-functional coordination.
- **LLM Post-Training:** GRPO, DPO, LoRA, Reward Modeling, Chain-of-Thought, SFT.
- **Evaluation & Benchmarking:** Benchmark Design, Rubric Design, Automated Evaluators, Failure Taxonomy.
- **Agent Infrastructure:** Multi-Agent Orchestration, Tool Use, MCP, Long-Horizon Agents
- **ML at Scale:** Distributed Training, MoE Architecture, Perception Data Flywheels, Modal, Synthetic Data Pipelines
- **Programming:** Python, C++, PyTorch, TensorFlow, Transformers, SQL, Cloud Platforms

## EDUCATION

---

- Stanford University, AI Graduate Program (2022-2024)
- Penn State University, B.S. Computer Science (2013-2017)