Michael Simon

mlsimon.com

Email : mlsimon@cs.ucla.edu

github.com/mlsimon734

Education

- **University of California, Los Angeles** — Los Angeles, CA
  Masters of Science in Computer Science — *Sept. 2024 – June 2026*
  - Advised by Eunice Jun
- **University of California, Los Angeles** — Los Angeles, CA
  Bachelors of Science in Computer Science — *Oct. 2020 – March 2024*

Experience

- **Excavaite Inc** — Palo Alto, CA
  CTO & Co-founder — *February 2024 – April 2026*
  - Built terabyte-scale patent corpus with RaBitQ-indexed vector search
  - Designed agentic systems for patent leakage risk prevention, automated patent drafting, and claim chart generation; instrumented automated evals to gate model and prompt changes
  - Raised pre-seed ($400k+); landed 6-figure pilot with Fortune 500 partner; managed 8+ contractor team
- **VAIL Lab** — Los Angeles, CA
  Undergraduate Researcher — *June. 2023 – June. 2024*
  - Worked on Diffusion Models for synthetic driving data generation, resulting in NeurIPS 2024 publication
  - Created unified terabyte scale driving dataset by combining 4k Youtube driving footage and open source datasets
- **Confluent** — Mountain View, CA
  Software Engineer Intern — *June. 2023 – Aug. 2023*
  - Created a last-resort metadata recovery tool (Java) for the Kafka Metadata Team to reduce metadata-related Kafka cluster failure recovery time from hours to minutes
  - Wrote a user-facing runbook and design document to support proper tool usage and ongoing development
- **Meta** — Menlo Park, CA
  Software Engineer Intern — *June 2022 – Aug. 2022*
  - Created a centralized Hack (PHP) Code Generation framework for the Metrics Platform Governance team; reduced time needed to onboard data into the platform from ∼1 hr to ∼5 min
  - Improved script performance by adding parallelism, reducing runtimes by 4x
- **Agrofocal Technologies** — San Jose, CA
  Machine Learning Intern — *June 2021 – Nov. 2023*
  - Implemented a mixed-precision quantization pipeline, improving inference pipeline throughput over 80% and decreasing latency 2x
  - Implemented state-of-the-art model architecture, improving object detection accuracy over 0.15 mAP

Projects

- **CLIP Explainability Study** | *PyTorch, NumPy* — March 2023
  - Studied CLIP (Contrastive Language-Image Pretraining) models across varying pre-training dataset sizes; visualized attention heatmaps and benchmarked image classification performance to characterize how data scale shapes learned representations

Selected Publications

* *denotes equal contribution*

- **Syrup: An Intermediate Representation for Bidirectional Human-AI Scientific Reasoning**. *Under review at the 39th ACM Symposium on User Interface Software and Technology (UIST 2026).* **Michael Simon***, Yuwei Xiao*, Naisha Agarwal, Eunice Jun.
- **SimGen: Simulator-conditioned Driving Scene Generation**. *In Proceedings of the 38th Advances in Neural Information Processing Systems (NeurIPS 2024).* Yunsong Zhou, **Michael Simon**, Zhenghao Peng, Sicheng Mo, Hongzi Zhu, Minyi Guo, Bolei Zhou.