

Eric Xue

https://www.ericxue.com

ex2156@columbia.edu

Mobile : (646) 679-0601

EDUCATION

Columbia University

Master of Science in Computer Science; GPA: 4.13/4.00

New York, NY

Sep 2025 - May 2026

University of Toronto

Honours Bachelor of Science in Computer Science; GPA: 3.93/4.00

Toronto, ON

Sep 2021 - Jul 2025

SKILLS

Languages: Python, C, TypeScript/JavaScript, Java, SQL

ML/AI: PyTorch, TensorFlow, ONNX, Hugging Face Transformers, vLLM, OpenAI Responses API, Weights & Biases, Slurm

Web & Infrastructure: Next.js, React.js, FastAPI, GCP (Cloud Run, GCE, Cloud Storage), PostgreSQL, Docker, Git

EXPERIENCE

DREAM Lab @ UIUC

Research Assistant

Remote

Jun 2022 - Mar 2025

- First-authored a research study on improving adversarial robustness of condensed datasets via curvature regularization, proposing a novel training objective that improved robust accuracy by up to **47%** over prior condensation methods on standard CV benchmarks. Published at **AAAI 2025**.
- Led development of AutoModel, a Python **multi-agent LLM** framework that orchestrates expert agents to automate end-to-end computer vision pipelines, validated across 3 standard CV benchmarks, matching human-tuned baselines.
- Managed experimentation infrastructure: wrote shell scripts to orchestrate **Slurm** batch jobs across 4 **A100** GPUs, automating experiment dispatch, output logging, and result aggregation over 100+ training configurations.

SocialAI Research Group @ UofT

Research Assistant

Toronto, ON

Sep 2022 - Sep 2024

- Led a full rewrite of Gem, a Python **multi-agent RL** framework, reducing the codebase by 50% while modularizing the environment and agent APIs to support plug-and-play experiment configuration for external research teams.
- Collaborated with 10+ researchers across Google DeepMind, Meta, and Vector Institute, incorporating cross-team feedback to shape Gem's API design and enable 5+ internal and external groups to build with the framework.

Cowork AI

Machine Learning Engineer

Toronto, ON

Sep 2023 - Dec 2023

- Fine-tuned **Stable Diffusion** model with **LoRA** on customer product image datasets, systematically evaluating training configurations to optimize for lighting fidelity, angle variation, and text legibility in generated ads.

aUToronto

Machine Learning Engineer

Toronto, ON

Sep 2022 - Apr 2023

- Trained a unified **YOLOv5** detector using **PyTorch Lightning** for autonomous vehicles, consolidating 3 previously separate models into a single multi-class detector, improving inference efficiency and detection accuracy by 10%.
- Built the next-generation **training dataset** for the team's perception stack: captured thousands of real-world images, defined the annotation schema for 5 object classes, and labeled 5,000+ bounding box annotations in **CVAT**.
- Integrated **Weights & Biases** to the training workflow, setting up experiment tracking and automated hyperparameter sweeps that replaced ad-hoc manual tuning and cut experiment iteration time by 60%.

SELECTED PROJECTS

Essayly

Full-Stack Engineer & Founder

New York, NY

Mar 2025 - Oct 2025

- Independently built and deployed a full-stack **LLM-powered** writing platform using **Next.js**, **Supabase PostgreSQL**, with **GitHub Actions** CI/CD and **Sentry** monitoring, serving 50+ users and generating \$1,000+ in revenue.
- Built data pipelines to prepare 1,000+ essay samples for fine-tuning LLMs, and implemented a multi-stage writing pipeline with **chunked embedding retrieval** for reference matching, reducing user rewrite attempts by **40%**.

PUBLICATIONS

Towards Adversarially Robust Condensed Dataset by Curvature Regularization

AAAI 2025

Learning to Imitate with Less: Efficient Individual Behavior Modeling in Chess

TMLR (2025)

Granular Analysis of Pretrained Object Detectors

ICAIIC 2022