

# Xiaole (Charlotte) Su

Boston, MA | (774) 329-5315

[su.xiaol@northeastern.edu](mailto:su.xiaol@northeastern.edu) | <https://xiaolesu.com> | [linkedin.com/in/xiaolesu/](https://www.linkedin.com/in/xiaolesu/) | [github.com/suxls](https://github.com/suxls)

## EDUCATION

### Northeastern University

Boston, MA

B.S. in Computer Science and Mathematics

Expected April 2027

**Courses:** Machine Learning I, Programming Languages, Algorithms, Linear Algebra, Graph Theory, Object Oriented Design

**Teaching Assistant:** Lab lead for Computer Fundamentals I

**Activity:** E-board at Rev (Startup Incubator), Software Developer at Sandbox (Software Consultancy), Volleyball, Dance

## EXPERIENCE

### Osmosis (YC W25)

January 2026 – Present

ML Engineer Co-op

San Francisco Bay Area

- Building a LLM **post-training** platform focused on reinforcement learning to align vertically specialized small AI models.
- Deploying a **Temporal**-orchestrated RL pipeline (**SkyPilot**, AWS/Lambda GPUs, SLIME/GSPO, Docker, FastAPI) for distributed training, automated data processing, and fault-tolerant checkpointing.
- Conducting independent research on **Lean 4** math autoformalization, iterating across **SFT** and **GRPO** training runs on **Qwen** models with Lean compiler feedback loops, with 2 small language models slated for open-source release.
- Curating multi-source datasets from public open-source corpora, generating synthetic **chain-of-thought** traces via **distillation** from frontier models like GLM-5 to bootstrap reasoning supervision for small model training.
- Running eval benchmarks against **open-source model baselines** to assess reasoning performance, guiding dataset curation and training iterations; co-authoring a company blog post documenting research findings and methodology.

### Bain & Company

January 2025 – June 2025

Platform Engineering Co-op

Boston

- Enhanced internal IaC platform, resolving account decommission errors and delivering detailed failure alerts.
- Administered **1000+** cloud accounts across Azure, AWS, and GCP, identifying and resolving automation failures.
- Created an AI pipeline that generates Terraform module code, reducing self-service IaC template rollout time by **96%**.
- Built Bain's first reusable **AWS Lambda** module, decreasing workload deployment time to <10 minutes for 20+ teams.
- Reduced AWS storage costs by **\$250K** annually by developing EC2 script to remove orphaned snapshots.

## PROJECT

### Gradually-Typed Programming Language Model

October 2025

- Designed and implemented a gradually-typed programming language with CESK abstract machine semantics in **OCaml**, modelled on academic research in Typed Racket and TypeScript.
- Built a complete language toolchain featuring an S-expression parser, static type validators, a linker, and a runtime with detailed state tracing and robust error diagnostics.
- Supported arithmetic computation, conditionals, loops, and mixed-type programs with both pre-runtime and runtime checks to enforce soundness across typed and untyped code.

### EyeCraft

September 2025

- **Ranked first** in entertainment track for **HackMIT** among 350+ projects created by over 1,000 undergraduate hackers.
- Developed a Minecraft mod using Fabric API, enabling handicap-accessible gameplay through alternative input.
- Integrated a live facial recognition pipeline using MediaPipe and **OpenCV**, mapping head poses, facial data, and voice commands to game controls; Trained neural networks on data with TensorFlow, achieving over 90% validation accuracy.

### KADA: Dance Formation Parser

September 2024

- **Earned 2nd place** in AI Track at Boston Hack in competition with 300+ hackers.
- Engineered an AI choreography tool leveraging DeepFace and OpenCV to generate top-down dance formation sequences.

## SKILLS

**Languages:** Python, Java, C#, OCaml, TypeScript, JavaScript, Go, Racket, HCL

**Tools:** Git, Bash, Docker, Kubernetes, HuggingFace, PyTorch, vLLM, SGLang, MLflow, Weights and Biases, OpenCV, Next.js

**Cloud & Infrastructure:** AWS, Lambda, Azure, GCP, SkyPilot, Temporal, Terraform Associate (Certified)