

SERAPH YANG

669-289-5716 | seraph.yang@stanford.edu | [linkedin.com/in/seraph-yang](https://www.linkedin.com/in/seraph-yang) | github.com/seraph-yang | seraphyang.com

EDUCATION

Stanford University

BS in Computer Science & Mathematics

Stanford, CA

Expected 2028

- **Relevant Coursework:** Deep Reinforcement Learning, Deep Learning for Computer Vision, Programming Abstractions, Probability for Computer Scientists, Mathematical Foundations of Computing, Linear Algebra, Differential Equations

EXPERIENCE

Software Development Engineer

Amazon

Jun. 2026 – Present

Sunnyvale, CA

- Modernized Amazon's internal **Kindle rendering infrastructure** by upgrading core rendering packages and migrating native UI components to **React Native Fabric**.
- Built **AI-assisted tooling** for automated **Git branch merging**, reducing manual integration effort and accelerating development workflows.
- Enhanced internal dialog-rendering frameworks, to support monorepo builds with **Bazel**.

Undergraduate Researcher

Stanford Artificial Intelligence Laboratory (SAIL)

Sep. 2025 – Present

Stanford, CA

- Proposed and implemented a latent-action world modeling framework based on JEPA for learning robot-relevant dynamics from **125K unlabeled videos** without action supervision.
- Investigated continuous versus discrete latent action representations, demonstrating superior transferability and achieving up to **1.76× scene-change sensitivity** across multiple regularization strategies.
- Developed an **MPC-based planning pipeline** leveraging learned latent dynamics to evaluate **512 candidate trajectories** for goal-conditioned robot control.
- Trained a latent-action controller on the DROID dataset using inverse-dynamics pseudo-labels, enabling video-to-robot transfer while maintaining **>97% rollout agreement** with an inverse-dynamics oracle.

PROJECTS

Abductive Event Inference using Temporal Graph RAG | Python, ChromaDB, Qwen

- Developed and deployed a Chroma DB vector database using Qwen3-0.6B embedding model to support temporal event retrieval and inference.
- Improved accuracy by 6.5% over baseline for submission to Task 12 of SemEval 2026, achieving 69% accuracy.

StudyAP | Next.js, Turso, Drizzle

- Built AI platform using **OpenAI & Deepseek API** generating AP practice problems, scaling to **1,450+ student users**.
- Implemented grading system for multiple-choice and free-response, outperforming existing prep tools.
- Pitched to VCs, securing mentorship and refining product vision.

OCR with OpenCV EAST and Tesseract | Python, OpenCV, PyTorch, Tesseract

- Developed **computer vision project** for text detection and recognition in images.
- Used pretrained **EAST neural network** to generate bounding boxes via confidence scores and geometry tensors.
- Applied **non-max suppression** to refine boxes, then extracted text with **Tesseract OCR**.
- Delivered accurate transcription from noisy inputs, showcasing **machine learning** in document analysis.

TECHNICAL SKILLS

Languages: Python, C/C++, Java

Frameworks: React, Next.js

Developer Tools: Git, Bazel, Claude Code

Libraries: NumPy, pandas, Matplotlib, PyTorch, TensorFlow, seaborn, RadGraph

INTERESTS

Golf, Squash, Charcoal sketching, Rock climbing, Chess