

Viren Khandal

virenkhandal.github.io | github.com/virenkhandal

Email: virenkhandal@berkeley.edu

Mobile: +1-510-386-6222

EXPERIENCE

- Tesla** Fremont, CA
 - Senior Machine Learning Engineer** Dec 2022 – Present
 - Leading ML research and development across 4 global factories, spanning vision-guided robotics, large-scale manufacturing defect detection, and core ML infrastructure for an internal computer vision training platform.
 - Investigating the open problem of factory-aware foundation models – how to learn visual representations that encode manufacturing domain priors and generalize across diverse production environments, inspection modalities, and defect taxonomies.
 - Formulated robust transfer and imitation learning approaches for adapting large pretrained vision models to novel factory domains with minimal labeled data, studying the conditions under which few-shot domain adaptation preserves detection fidelity.
 - Researched knowledge distillation strategies for compressing high-capacity teacher networks into lightweight student models deployable at real-time edge inference latencies, analyzing accuracy-efficiency trade-offs in resource-constrained manufacturing settings.
 - Designed and evaluated production computer vision systems for automated quality inspection, contributing empirical insights on deploying learned models in high-throughput, low-intervention industrial pipelines.
- Ponder (Stealth)** Remote
 - Co-Founder** Oct 2024 – Jan 2026
 - Led full backend and ML development for an AI app that generates multi-perspective audio podcasts from any user query, with each perspective grounded in cited research and credible sources.
 - Architected retrieval-augmented generation pipelines to synthesize web-scale information into structured, multi-viewpoint audio content with source attribution.
 - Built end-to-end backend infrastructure including API design, audio generation, and serving systems to support real-time podcast creation at scale.
- University of California, Berkeley – VeHICal** Berkeley, CA
 - Research Engineer** May 2021 – Dec 2022
 - Led research on safe autonomous-to-human perception handoffs under Professors Sanjit Seshia and Bjoern Hartmann (vehical.org), formalizing stochastic models for verified human-autonomy interaction.
 - Developed VR-based human-in-the-loop experimental tools to study driver takeover scenarios with high ecological validity.
- Stanford Autonomous Systems Lab** Stanford, CA
 - Visiting Researcher** Apr 2022 – Dec 2022
 - Investigated out-of-distribution (OOD) events in human-autonomy interaction under Professor Marco Pavone (stanfordasl.github.io), developing frameworks for safe real-time control handoffs during OOD scenarios.

EDUCATION

- University of California, Berkeley** Berkeley, CA
 - B.A. Computer Science & Applied Mathematics (Double Major)** Aug 2019 – May 2022
 - Specialization in Machine Learning and Artificial Intelligence
 - Coursework: Deep Learning, Machine Learning, Artificial Intelligence, Efficient Algorithms, Data Structures, Statistical Prediction, Numerical/Real/Complex Analysis, Advanced Linear Algebra, Abstract Algebra
- Stanford University** Palo Alto, CA
 - Certificate in Computer Science** May 2018 – Sept 2018
 - Coursework: Data Mining and Analysis, Client-Side Internet Technologies

TECHNICAL SKILLS

Languages: Python, C/C++, Go, Java, JavaScript, SQL

ML/AI: PyTorch, TensorFlow, OpenCV, Reinforcement Learning, LLMs, RAG

Infrastructure: AWS, GCP, Docker, Kubernetes