



ALVIN ZHANG

+41 077 265 39 87 ◊ alvn.zng@gmail.com ◊ <https://alvinzz.github.io/>

OBJECTIVES

I enjoy progressing practical projects by applying my expertise with state-of-the-art neural networks.

EDUCATION

EPFL Sept 2023 - Aug 2024

Ph.D., Computer and Communication Sciences

University of California, Berkeley 2016 - 2020

B.Sc., Electrical Engineering and Computer Science Major GPA: 3.875

SELECTED COURSEWORK:

Computer Vision (A+), Computer Graphics (A+), Robotic Manipulation and Interaction (A+), Machine Learning (A+), Artificial Intelligence (A+), Statistical Learning Theory (A), Feedback Control Systems (A)

INDUSTRY

Research Engineer May 2021 - May 2023

Matician (Autonomous Household Robotics Startup) *Mountain View, CA*

- Inaugural member of the Mapping team; lead and mentor future members; coordinate with the CEO and others.
- Research, integrate, and deploy real-time a centimeter-resolution 3-D mapping algorithm on a real-world robot.
 - Lead research and development of the central node in the autonomy stack.
 - Create visualizations and define the user-facing interface.
 - Fuse multi-modal inputs via a novel probabilistic interpretation inspired by the signed distance function; apply Bayesian inference using hardware-aware approximations to achieve centimeter-level resolution.
 - Probe timing and scaling of DSL operations on the chip; coordinate with the chip manufacturer to expose bugs and write tests.
- Identify unexpected anomalies during operation; incorporate this into internal state and use it to adapt behavior.
- Integrate IMU measurements into a Visual-SLAM system. Analyze sensor noise, determine root causes, and develop mitigation strategies to run on-bot.

Perception Software Engineer Jan 2021 - Apr 2021

Nuro (Autonomous Delivery Vehicle Startup) *Mountain View, CA*

- Review state-of-the-art literature on optical flow; train and evaluate industrial-sized models for optical flow.

Perception Engineer May 2020 - Jan 2021

Ike Robotics (Autonomous Trucking Startup), acquired by Nuro Jan 2021 *San Francisco, CA*

- Research, develop, and evaluate neural approaches for 3-D bounding box segmentation from LiDAR point clouds.
- Co-author, with Lance Martin, a blog post on [“Perception for Automated Trucking”](#).

RESEARCH

Semester Project Jan 2024 - Aug 2024

CLAIRE, EPFL with *Prof. Caglar Gulchere.* *Lausanne, CH*

- MULTIPLICATION: A CASE STUDY IN ALGORITHMIC EXECUTION IN NEURAL NETWORKS
- Explore trade-offs between LLM model architecture and training method on algorithmic execution tasks, including nearest-neighbor retrieval, counting, parity, and multiplicative distribution.

Semester Project Aug 2023 - Dec 2023

VILAB, EPFL with *Prof. Amir Zamir.* *Lausanne, CH*

- WHAT'S THE PROBLEM? LEARNING WHAT TO SOLVE AT TEST-TIME
- Investigate an energy-based model which, inspired by diffusion models, matches the gradient of the energy to a denoising target. Demonstrate a trade-off between inference time and output quality.

Research Intern

May 2022 - Aug 2022

Redwood Center for Theoretical Neuroscience, UC Berkeley *with Prof. Bruno Olshausen.*

Berkeley, CA

- Empirically demonstrate that a hierarchical generative model performs “explaining away” of nuisance factors; show that this process follows the theoretical neural dynamics of predictive coding and Bayesian inference.
- Extend NeRF by learning 3-D structural priors in a multi-scale volumetric rendering system, with the goal of improving reconstruction quality in sparse-view scenarios.

PUBLICATIONS

Generalized Skill Learning, Safety, and Exploration with Flow-Based Models [\[link\]](#)

May 2019

Workshop, Task-Agnostic Reinforcement Learning, ICLR

New Orleans, LA

- **Single-author paper.** Demonstrate a robotic Ant teaching itself to perform jumps and flips via unsupervised exploration and self-supervision: goal-conditioned RL with time between states as the cost function.
- Introduce a variant which promotes risk-averse exploration in safety-critical environments.

PROJECTS

A Geometric Introduction to Lie Theory [\[link\]](#)


Nov 2022 -

Tutorial - WIP

- Provide an intuitive, visual introduction to Lie theory through the motivating example of rigid-body motion.
- Introduce a novel formula for computing the logarithm of SE(3) that is faster and more numerically stable than existing methods.

Weathering with Moo: Simulating Building and Statue Weathering [\[link\]](#)

Spring 2020



Computer Graphics Final Project

- Simulate natural weathering and erosion of myriad materials by stochastic deformations of mesh surfaces.

Fun with Homographies [\[link\]](#)

Spring 2019



Computer Vision Final Project

- Create 3D visual illusions of wire sculptures that create different shapes following the perceiver's viewpoint.

Drawyer [\[link\]](#)

Spring 2018



Robotics Final Project

- Program a Sawyer robot to draw line sketches of reference pictures.

LIBRARIES

Named Tensors [\[link\]](#)

Nov 2021

PyTorch Library

- Support named dimensions for tensors. Reduce code ambiguity; enable better debugging and run-time checks. At the original time of its development, this provided more functionality than PyTorch's native implementation.

Type-Safe Coordinate Transforms [\[link\]](#)

Oct 2022

Rust Library

- Implement type-safety for transforms between coordinate systems: projecting a point in coordinate frame A into camera B raises a compile-time error; composing $G_a^b(t = 1)$ with $G_b^c(t = 2)$ raises a run-time error.

SKILLS

Python, Rust, C++, pandas, PyTorch, TensorFlow, CUDA, scikit-learn, OpenCV, numpy, Google Cloud / GCP, AWS, Docker, Kubernetes, Grafana, Large Language Models / LLMs, GPT, Llama, Machine Learning, Computer Vision, Deep Reinforcement Learning / Deep RL, Git / GitHub, OpenSCAD