

Utkarsh Singhal

Ph.D. Candidate @ UC Berkeley | s.utkarsh@berkeley.edu | utkarsh.ai | [Google Scholar](https://scholar.google.com/citations?user=utkarsh) | [GitHub](https://github.com/utkarsh)

About Me

Currently working on adding better physical dynamics into **robotic world models**.

More generally, my research explores how to learn **scalable** data-driven priors for **embodied AI**.

Education

University of California, Berkeley Dec 2025 (expected)
Ph.D., Electrical Engineering and Computer Science (EECS)
Thesis: *Scalable Learning and Integration of Physical Constraints for Embodied AI* (with Prof. Stella Yu)

University of California, Berkeley May 2018
B.S., Electrical Engineering and Computer Science (EECS) GPA: 3.9/4.0
B.A., Applied Mathematics (dual degree)

Selected publications

(Ongoing) **Physically plausible world models by combining video diffusion and latent space models**

[Utkarsh Singhal](#), [Stella X. Yu](#)

Description: Combining action-conditioned video diffusion with latent space models (VJEPa-2-AC) and regularization (REPA).

Test-time canonicalization by foundation models for robust perception: ICML'25

[Utkarsh Singhal*](#), [Ryan Feng*](#), [Atul Prakash](#), [Stella X. Yu](#)

Description: Test-time optimization using CLIP/SD priors makes embodied AI more robust to natural transforms.

How to guess a gradient: OPT Workshop @ NeurIPS'23

[Utkarsh Singhal*](#), [Brian Cheung*](#), [Kartik Chandra](#), [Jonathan Ragan-Kelley](#), [Joshua Tenenbaum](#), [Tomaso Poggio](#), [Stella X. Yu](#)

Description: Used NN structure to improve forward-mode gradients by 100 \times , explored training with second-order optimizers.

Learning to transform for generalizable instance-wise invariance: ICCV'23

[Utkarsh Singhal](#), [Carlos Esteves](#), [Ameesh Makadia](#), [Stella X. Yu](#)

Multi-spectral image classification with ultra-lean complex-valued models: HADR Workshop @ NeurIPS'22

[Utkarsh Singhal](#), [Stella X. Yu](#), [Zackery Steck](#), [Scott Kangas](#), [Aaron Reite](#)

Fourier features let networks learn high-frequency functions in low-dimensional domains: NeurIPS'20

[Matt Tancik](#), [Pratul Srinivasan](#), [Ben Mildenhall](#), [Sara Fridovich-Keil](#), [Nithin Raghavan](#), [Utkarsh Singhal](#), [Ravi Ramamoorthi](#), [Jon Barron](#), [Ren Ng](#)

Wireless recording in the peripheral nervous system with ultrasonic neural dust: Neuron, July 2016

[Dongjin Seo](#), [Ryan M Neely](#), [Konlin Shen](#), [Utkarsh Singhal](#), [Elad Alon](#), [Jan M Rabaey](#), [Jose M Carmena](#), [Michel M Maharbiz](#)

Previous experience

Research Intern, Adobe (Host: Michael Gharbi, Connelly Barnes) May 2019 – Jan 2020

- Discovered an optimization for PatchMatch (a content-aware fill algorithm), accelerating it by 20 \times . ([Patent](#))
- Developed a *differentiable* PatchMatch and its custom memory-efficient autograd code in PyTorch.

Oz Vision Project, UC Berkeley (Prof. Ren Ng) August 2018 – May 2019

- Worked on a team building a millisecond-resolution system to track and stimulate retina with a laser. ([Patent](#))
- Wrote low-latency networking code to synchronize rendering across machines at millisecond speed.

Skills

Computer vision, diffusion models, control theory, reinforcement learning, PyTorch, SLURM, Lightning, Docker, W&B, Hydra, DDP/FSDP, LoRA, computer vision, vision-language action models, transformers, gaussian splatting, IsaacLab, sim2real, imitation learning, signal processing, electronics, 3D printing/prototyping.