Rayhan Zirvi

Generative Modeling · California Institute of Technology

rayhanzirvi@gmail.com
rayhanzirvi.com
Google Scholar

EDUCATION

California Institute of Technology — Pasadena, CA
B.S., Computer Science — *2023 – 2027 (expected)*
GPA: 4.2/4.0
Selected coursework: Advanced Topics in Machine Learning[G], Analysis & Design of Algorithms[G], Markov Chains & Discrete Stochastic Processes[G], Scientific Machine Learning[G], Combinatorial Analysis[G], Probability Models, Statistical Inference, Game Theory ([G] = graduate course)

RESEARCH EXPERIENCE

Guibas Lab, Stanford University — Stanford, CA
Research Intern, Advised by Prof. Leonidas Guibas — *Mar 2026 – Sep 2026*

- Co-led AnisoFlow (co-first author): replaces hand-designed prediction targets for flow models with a fully learnable one, improving JiT from 1.90 to 1.55 FID on ImageNet 256×256 with the architecture, loss, and sampler unchanged
- Contributed to Decoding Without a Decoder, a post-training method that lets diffusion transformers predict final pixels directly, eliminating the VAE decoder from image, text-to-image, and video models

Anima AI + Science Lab, Caltech — Pasadena, CA
Undergraduate Researcher, Advised by Prof. Anima Anandkumar and Prof. Bahareh Tolooshams — *Feb 2024 – May 2026*

- Co-led DiffStateGrad (co-first author, ICLR 2025): projects measurement gradients onto a low-rank subspace of the diffusion state, improving the reconstruction quality and robustness of diffusion-based inverse problem solvers
- Co-led EquiReg (co-first author, TMLR 2026): introduces equivariance-based regularization that discourages off-manifold sampling trajectories, improving reconstructions across imaging and PDE inverse problems

PUBLICATIONS

* denotes equal contribution

Learning Anisotropic Prediction Targets for Flow Models
Rayhan Zirvi*, Hansheng Chen*, Gordon Wetzstein, Leonidas Guibas
arXiv preprint, 2026 · arXiv

Decoding Without a Decoder: Direct Pixel Prediction from Diffusion Transformers
Hansheng Chen, Shengqu Cai, Jan Ackermann, **Rayhan Zirvi**, Gordon Wetzstein, Xinchen Yan, Leonidas Guibas
Under review, 2026

EquiReg: Equivariance Regularized Diffusion for Inverse Problems
Bahareh Tolooshams*, Aditi Chandrashekar*, **Rayhan Zirvi***, Abbas Mammadov, Jiachen Yao, Chuwei Wang, Anima Anandkumar
Transactions on Machine Learning Research (TMLR), 2026 · arXiv · Code

Diffusion State-Guided Projected Gradient for Inverse Problems
Rayhan Zirvi*, Bahareh Tolooshams*, Anima Anandkumar
International Conference on Learning Representations (ICLR), 2025 · arXiv · Code

AWARDS

2026 Barry Goldwater Scholarship
2026 Henry Ford II Scholar Award (best academic record in computer science at Caltech)
2026 Taylor Lawrence Summer Undergraduate Research Fellowship
2025 George W. Housner Discovery Fund Award
2024 Caltech Summer Undergraduate Research Fellowship

INDUSTRY EXPERIENCE

Meta Menlo Park, CA
Software Engineer Intern, Machine Learning *Jun 2025 – Sep 2025*

- Designed a multimodal LLM pipeline to cluster image, video, and text content and describe each for safety review
- Built batch-inference infrastructure to run the pipeline daily on thousands of items, informing other teams' content classification models

COMPETITIVE PROGRAMMING

ICPC North America Championship, Honors — Southern California Regional Champion
Codeforces Master — top 40 in the US
Meta Hacker Cup Round 3 Qualifier (Semifinalist) — top 1% internationally
USA Computing Olympiad, Platinum Division — top 200 in the US
Google Kick Start — 11th in the US

TALKS & POSTER PRESENTATIONS

Learning Anisotropic Prediction Targets for Flow Models
Rayhan Zirvi, Hansheng Chen, Gordon Wetzstein, Leonidas Guibas
Talk, SURF Seminar at Caltech, 2026

Diffusion State-Guided Projected Gradient for Inverse Problems
Rayhan Zirvi, Bahareh Tolooshams, Anima Anandkumar
Poster, International Conference on Learning Representations (ICLR), Singapore, 2025

Diffusion State-Guided Projected Gradient for Inverse Problems
Rayhan Zirvi, Bahareh Tolooshams, Anima Anandkumar
Talk, SURF Seminar at Caltech, 2024

SERVICE

2027 Reviewer, International Conference on Learning Representations (ICLR)