Karhan Kayan

Ph.D. Candidate in Computer Science, Princeton University
karhankayan.com | karhan@princeton.edu | linkedin.com/in/karhankayan | github.com/karhankayan | Google Scholar

Research Interests

3D computer vision, procedural generation and synthetic data, robot learning and simulation, camera calibration and pose estimation, benchmarks for foundation models. 6 publications at CVPR, ICCV, ECCV, NeurIPS, and 3DV.

Education

Princeton University — Princeton, NJ
Ph.D. in Computer Science (expected 2027) — *Sept. 2022 – Present*

- Advisor: Jia Deng, Princeton Vision & Learning Lab. GPA: 3.95/4.00
- Coursework: Advanced Computer Vision, Robot Perception and Learning, Advanced Algorithm Design, Computational Models of Cognition, Computational Imaging, Advanced Computer Networks

University of Chicago — Chicago, IL
Honors B.S. in Computer Science and Honors B.S. in Mathematics — *Sept. 2018 – June 2022*

- Only student in the Class of 2022 to earn two Honors B.S. degrees. GPA: 3.77/4.00 (CS: 3.87/4.00)
- Graduate and honors coursework: Honors Analysis, Honors Algorithms, Honors Abstract Algebra, Machine Learning, Deep Learning, Theory of Machine Learning, Computer Vision, Computer Graphics, Robotics, Operating Systems, Information Theory, Topological Data Analysis, Measure Theory, Complex Analysis, Stochastic Processes, Partial Differential Equations, Topology

Research Experience

Princeton Vision & Learning Lab — Princeton, NJ
Graduate Research Assistant (Advisor: Jia Deng) — *July 2022 – Present*

- **Infinigen** — Core developer of Infinigen (CVPR 2023, 2024), a procedural generator of unlimited photorealistic 3D training data for vision and robotics, built entirely from randomized mathematical rules with automatic ground truth for depth, surface normals, optical flow, and segmentation. **7,100+ GitHub stars**; featured in NVIDIA Isaac Sim's official documentation for synthetic data generation. Co-led Infinigen Indoors (CVPR 2024) and built the fluid simulation system for Infinigen Nature (CVPR 2023).
- **ProcFunc** — Developed the ground-truth annotation system and procedural materials for ProcFunc (under review), the next-generation Infinigen framework for procedural 3D generation in Python.
- **Princeton365** — Co-led Princeton365 (ICCV 2025), a benchmark of 365 diverse indoor, outdoor, and object-scanning videos (2.1M frames) with millimeter-accurate camera-pose ground truth and synchronized stereo RGB and IMU, collected via a novel 360°-camera and calibration-board pipeline. Developed BundlePnP, a Ceres-based nonlinear least-squares method that jointly optimizes camera and calibration-board poses under fixed board geometry, achieving 2.88 mm mean ATE validated against Vicon motion capture (vs. 11.4 mm for COLMAP). Introduced a scale-aware, optical-flow-based error metric that enables cross-scene comparison, addressing a fundamental limitation of ATE.
- **UniQA-3D** — Co-led UniQA-3D (3DV 2025), a VQA-format benchmark comparing VLMs, specialized 3D models, and human subjects on core 3D vision tasks: state-of-the-art VLMs fail at basic 3D tasks, specialized models are accurate but brittle under geometric perturbations, and Vision Transformers align more closely with human 3D perception than CNNs.
- **DARPA program** — Led the student team on a $1M DARPA program on foundation models for materials synthesis: coordinated the participating Ph.D. students, authored program reports, and built tools for automated SEM imaging of material microstructure, automated discovery of physical invariances via symbolic regression, and a tool-augmented LLM agent for synthesis-recipe discovery.
- **Robot learning** — Building a robotics simulation toolkit that procedurally generates scenes and expert trajectories using the NVIDIA Newton physics engine and Blender; trained imitation-learning policies for liquid handling on an xArm robot.
- **Infinigen-Articulations** — Designed procedural articulated 3D assets in Blender for Infinigen-Articulations (under review), a toolkit of generators for 18 articulated object categories with a physics-aware export pipeline to robotics simulators, shown to improve RL generalization and sim-to-real transfer.

- **InFlux** — Engineered the RTK positioning module (Septentrio Mosaic-X5) of the drone-based calibration system behind InFlux (NeurIPS 2025), the first large-scale benchmark for dynamic camera intrinsics (386 videos, 143K+ annotated frames), achieving centimeter-accurate positioning ground truth.
- **Depth from defocus** — Performed the camera calibration across the full focus stack of a DSLR for the real-world benchmark of Zero-shot Depth from Defocus (ECCV 2026, **oral presentation**), and assisted with benchmark filming.
- **DancingThings** — Proved two approximation theorems underpinning DancingThings (under review), the first optical flow dataset focused on general non-rigid motion, and implemented its ground-truth system with realistic, fast rendering.

Booth School of Business, University of Chicago — Chicago, IL
Statistics/ML Research Assistant (Advisor: Bryon Aragam) — *June 2021 – Jan. 2023*

- Researched graphical models, causal inference, and deep generative models; presented findings at the UChicago Center for Data and Computing.
- Implemented Variational Autoencoder architectures in PyTorch to study causal inference under latent confounding and the statistical identifiability of such methods.

Paparrizos Lab, University of Chicago — Chicago, IL
Machine Learning Research Assistant (Advisors: John Paparrizos, Michael Franklin) — *July 2020 – Dec. 2021*

- Implemented and released `grailts`, a Python package for GRAIL (VLDB 2019) time-series representation learning: kernel-based similarity (SINK), landmark-based low-rank representations, and matrix sketching.
- Implemented fast approximate k-nearest-neighbor search over the learned representations using product quantization (PQ/OPQ), preserving similarity measures beyond Euclidean distance.
- Developed and evaluated a method for scaling Granger-causality inference to large time-series collections: pruning the quadratic pairwise search with approximate k-NN achieves linear scaling with comparable accuracy in experiments on 128 UCR datasets and semi-synthetic benchmarks up to 1M series; designed the ARIMA-based benchmark generator with known causal ground truth (unpublished manuscript).

Mathematics REU, University of Chicago — Chicago, IL
Research Participant (Advisors: Minh-Tam Trinh, Peter May) — *June 2019 – Sept. 2019*

- Studied modular forms and analytic number theory; authored an expository paper on the orthogonal decomposition theorem for modular forms, drawing on algebraic geometry, Fourier analysis, and group theory.

Industry Experience

NOVAmedtek — Ankara, Turkey (Remote)
Computer Vision Intern — *Sept. 2021 – Oct. 2021*

- Wrote the image-processing software for an infrared vein visualization system in Python with OpenCV, implementing a published research method as a working prototype.
- Built the end-to-end pipeline using histogram equalization, adaptive thresholding, morphological transformations, and contour approximation.

Peer-Reviewed Publications

1. [ECCV 2026] Yiming Zuo*, Hongyu Wen*, Venkat Subramanian*, Patrick Chen, **Karhan Kayan**, Mario Bijelic, Felix Heide, Jia Deng. Zero-shot Depth from Defocus. *European Conference on Computer Vision (ECCV)*, 2026. *Equal contribution. **Oral presentation**

2. [NeurIPS 2025] Erich Liang, Roma Bhattacharjee, Sreemanti Dey, Rafael Moschopoulos, Caitlin Wang, Michel Liao, Grace Tan, Andrew Wang, **Karhan Kayan**, Stamatis Alexandropoulos, Jia Deng. InFlux: A Benchmark for Self-Calibration of Dynamic Intrinsics of Video Cameras. *Neural Information Processing Systems (NeurIPS)*, 2025.

3. [ICCV 2025] **Karhan Kayan***, Stamatis Alexandropoulos*, Rishabh Jain, Yiming Zuo, Erich Liang, Jia Deng. Princeton365: A Diverse Dataset with Accurate Camera Pose. *International Conference on Computer Vision (ICCV)*, 2025. *Equal contribution

4. [3DV 2025] **Karhan Kayan***, Yiming Zuo*, Maggie Wang, Kevin Jeon, Jia Deng, Thomas L. Griffiths. Towards Foundation Models for 3D Vision: How Close Are We? *International Conference on 3D Vision (3DV)*, 2025. *Equal contribution

5. [CVPR 2024] **Karhan Kayan***, Alexander Raistrick*, Lingjie Mei*, David Yan, Yiming Zuo, Beining Han, Hongyu Wen, Meenal Parakh, Stamatis Alexandropoulos, Lahav Lipson, Zeyu Ma, Jia Deng. Infinigen Indoors: Photorealistic Indoor Scenes using Procedural Generation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. *Equal contribution. **7,100+ GitHub stars**

6. [CVPR 2023] Alexander Raistrick*, Lahav Lipson*, Zeyu Ma*, Lingjie Mei, Mingzhe Wang, Yiming Zuo, **Karhan Kayan**, Hongyu Wen, Beining Han, Yihan Wang, Alejandro Newell, Hei Law, Ankit Goyal, Kaiyu Yang, Jia Deng. Infinite Photorealistic Worlds using Procedural Generation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. *Equal contribution. **7,100+ GitHub stars**

Preprints and Papers Under Review

1. Alexander Raistrick, **Karhan Kayan**, Jack Nugent, David Yan, Lingjie Mei, Meenal Parakh, Hongyu Wen, Dylan Li, Yiming Zuo, Erich Liang, Jia Deng. ProcFunc: Function-Oriented Abstractions for Procedural 3D Generation in Python (next-generation Infinigen framework). arXiv:2604.26943. Under review.

2. Stamatis Alexandropoulos, **Karhan Kayan**, Jia Deng. *DancingThings: Synthetic Data for Optical Flow with Spline-based Procedural Deformation.* Under review.

3. Abhishek Joshi, Beining Han, Jack Nugent, Max Gonzalez Saez-Diez, Yiming Zuo, Jonathan Liu, Hongyu Wen, Stamatis Alexandropoulos, **Karhan Kayan**, Anna Calveri, Tao Sun, Gaowen Liu, Yi Shao, Alexander Raistrick, Jia Deng. Procedural Generation of Articulated Simulation-Ready Assets. arXiv:2505.10755. Under review.

4. **Karhan Kayan**. Orthogonal Decomposition of Modular Forms. University of Chicago Mathematics REU, 2019.

Honors and Awards

Only student in the University of Chicago Class of 2022 to earn two Honors B.S. degrees (2022)

Sigma Xi, Full Member (elected 2022)

Neubauer International Scholarship, University of Chicago (2018–2022)

Quad Undergraduate Research Grant, University of Chicago (2×, 2020–2021)

SEAS Travel Grant, Princeton University (2×)

Dean's List, University of Chicago (2018–2019)

Turkish national team, 25th Balkan Olympiad in Informatics (2017)

4th, 5th, and 6th place, Turkish National Olympiad in Informatics rounds (2017)

Silver Medal, Istanbul Science Olympiad in Informatics (2017)

Ranked 44th among 1.11 million students, Turkish National High School Entrance Exam (2013)

Invited Talks

Integrating Generative Models and Causal Inference with Applications in Fair Machine Learning. Center for Data and Computing, University of Chicago, Aug. 2021

Teaching

Princeton University — Princeton, NJ
Teaching Assistant, COS 529 — Advanced Computer Vision (graduate) — *Spring 2025, Fall 2025*

Professional Service

- Maintainer, ImageNet (Jan. 2026 – Present)
- Reviewer: CVPR (2025–2026), ICCV (2025), ECCV (2026), NeurIPS (2026), BMVC (2026)
- Olympiad and Programming Mentor, Bahcesehir High School for Science and Technology (2015–2018): lessons on C++, algorithms, and discrete mathematics for national informatics olympiad students

Mentorship

- Stamatis Alexandropoulos, *Ph.D. student, Princeton* 2023–2024
- Anlon Zhu, *Undergraduate, Princeton → Software Engineer, Stripe* 2024–2025
- Maggie Wang, *Undergraduate, Princeton → Sierra AI* 2024
- Rishabh Jain, *Undergraduate, Princeton* 2024–2025
- Dylan Li, *Undergraduate, Princeton* 2024–2025
- Kevin Jeon, *Undergraduate, Princeton* 2024–2025
- Erik Dyer, *Undergraduate, Princeton* 2024–Present
- Tian Pu, *Undergraduate, Princeton* 2025–Present
- Sai Minnal, *Undergraduate, University of Pennsylvania* 2026–Present
- Junze Xie, *Undergraduate, Princeton* 2026–Present

Technical Skills

Languages: Python, C/C++, CUDA, SQL, R, MATLAB, Rust (basic), LaTeX
ML and Vision: PyTorch, OpenCV, COLMAP, Ceres, ROS, OpenGL, Trimesh, NumPy, SciPy, pandas, scikit-learn
Simulation and Tools: Blender, NVIDIA Isaac Sim/Lab, NVIDIA Newton, Docker, Linux, Git, Slurm
Methods: camera calibration and multi-view geometry, nonlinear least-squares optimization and bundle adjustment, procedural generation, imitation learning, SLAM and 3D benchmarking, approximate nearest-neighbor search, causal inference