Declan Kutscher

declank@berkeley.edu | https://www.linkedin.com/in/detk | https://d3tk.github.io

EDUCATION

University of Pittsburgh, *M.S.* in Computer Science **January 2024 - December 2025**
Thesis: *On the Effectiveness of Pretrained Models For Remote Sensing*

University of Pittsburgh, *B.S.* in Computer Science, *Magna Cum Laude* **August 2021 - April 2024**

RESEARCH EXPERIENCE

Research Engineer (BAIR, Darrell Group) **December 2024 - Present**
University of California, Berkeley

- Develop **score-based generative models** for SAR **blind inverse problems** from adaptively selected subapertures without degrading downstream performance
- Initiate **petabyte-scale self-supervised pretraining** of geospatial foundation models with the goal of advancing geospatial visual reasoning performance
- ***REOrdering Patches Improves Vision Models*** **(NeurIPS 2025):** Maintained large-scale codebase and training infra for long-sequence vision tasks
- Boosted ImageNet-1K (**+3.0%**) and FMoW (**+13%**) via **reinforcement-learned** patch reordering

Visiting Student Researcher (NeuroAgents Lab) **June 2025 - Present**
Carnegie Mellon University

- Develop a **reinforcement learning** framework leveraging **world models** to study **self-supervised skill discovery** and **intrinsic motivation**
- Investigate the **emergence of motor primitives** as a computational basis of autonomy in animals (fruit flies)

Research Assistant (Dr. Xiaowei Jia) **October 2022 - December 2025**
University of Pittsburgh

- ***Physics-Guided Fair Graph Sampling*** **(AAAI 2025):** Integrated PDE-based heat transfer into graph models to reduce bias in water temperature prediction; engineered fairness-relevant features
- ***Bringing Vision to Climate*****:** Built CNN–LSTM model (**PyTorch**) predicting stream depth from 18K hourly images and climate data; cut error by 40% with 25% of training data, improving transferability

PUBLICATIONS

D. Kutscher, "*On the Effectiveness of Pretrained Models for Remote Sensing*," M.S. thesis, Dept. of Computer Science, School of Computing and Information, Univ. of Pittsburgh, Pittsburgh, PA, 2025, forthcoming.

D. Kutscher, D. M. Chan, Y. Bai, T. Darrell, R. Gupta. *"REOrdering Patches Improves Vision Models"*. 2025. Neural Information Processing Systems (NeurIPS) 2025. [Online]. Available: https://neurips.cc/virtual/2025/loc/san-diego/poster/116773

He, E., **Kutscher, D.**, Xie, Y., Zwart, J., Jiang, Z., Yao, H., and Jia, X. (2025). Physics-Guided Fair Graph Sampling for Water Temperature Prediction in River Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, *39*(27), 28070-28078. https://doi.org/10.1609/aaai.v39i27.35025

D. Kutscher, A. Whisnant, "Leveraging LLMs for Data Coding," *Carnegie Mellon University, Software Engineering Institute's Digital Library*. Software Engineering Institute, White Paper, 04-Nov-2024 [Online]. Available: https://insights.sei.cmu.edu/documents/6210/DM24-1560.pdf

INDUSTRY EXPERIENCE

CERT Risk & Resilience Intern — May 2024 - August 2024

Carnegie Mellon University, Software Engineering Institute

- Implemented LLM-based data labeling for the insider threat database (1,800 cases, 5-8 documents per case), significantly reducing manual processing time using the OSS Ollama
- Integrated RAG to provide contextual awareness, enabling interactive labeling and improving extraction accuracy for IIDES schema objects
- Optimized performance through targeted prompt engineering and JSON segmentation strategies for complex structured outputs

Machine Learning Research Intern — June 2023 - August 2023

ThayerMahan, Inc. Groton, CT.

- Fine-tuned industrial manufacturing EfficientAD model for patch-level anomaly detection, integrating post-processing techniques to enhance segmentation quality and saliency
- Improved end-to-end pipeline for object detection reporting in gigapixel SAS imagery: patching, anomaly detection, mask stitching, and candidate object extraction
- Developed an automated candidate object extraction method, improving downstream classification accuracy by 2% without adding runtime overhead

PROJECTS

Intel OpenVino Toolkit Anomalib - *Contributor*

- Library for benchmarking, developing, and deploying deep learning anomaly detection algorithms in manufacturing
- Implemented, debugged, and tested the EfficientAD model support for the library

AWARDS

- NeurIPS Scholar Award, Conference on Neural Information Processing Systems (NeurIPS), 2025