

Jacob Mitchell Springer

✉ jspringer@cmu.edu

📄 Jacob Mitchell Springer

🌐 jakespringer

📧 sprin.xyz

Education

Carnegie Mellon University

Ph.D., Machine Learning

Pittsburgh, PA

August 2022 – Present

Swarthmore College

B.A., Mathematics and Computer Science

Swarthmore, PA

August 2017 – May 2022

Publications & Manuscripts

<https://scholar.google.com/citations?user=niZiN38AAAAJ>

- **Springer, Jacob M.**; Kotha, Suhas; Fried, Daniel; Neubig, Graham; Raghunathan, Aditi. Repetition Improves Language Model Embeddings. Preprint.
- **Springer, Jacob M.**; Nagarajan, Vaishnavh; Raghunathan, Aditi. Sharpness-Aware Minimization Enhances Feature Learning. ICLR 2024.
- Kotha, Suhas.; **Springer, Jacob M.**; Raghunathan, Aditi. (2023). Understanding Catastrophic Forgetting in Language Models via Implicit Inference. ICLR 2024.
- Jones, Haydn T; **Springer, Jacob M**; Kenyon, Garrett T; Moore, Juston. If You've Trained One You've Trained Them All: Inter-Architecture Similarity Increases With Robustness. Conference on Uncertainty in Artificial Intelligence. 2022. (Oral)
- **Springer, Jacob M**; Mitchell, Melanie; Kenyon, Garrett T. A Little Robustness Goes a Long Way: Leveraging Robust Features for Targeted Transfer Attacks. Advances in Neural Information Processing Systems (NeurIPS). 2021.
- **Springer, Jacob M**; Mitchell, Melanie; Kenyon, Garrett T. Uncovering Universal Features: How Adversarial Training Improves Adversarial Transferability. ICML 2021 Workshop on Adversarial Machine Learning. 2021. (Shorter version of above.)
- **Springer, Jacob M**; Mitchell, Melanie; Kenyon, Garrett T. Adversarial Perturbations Are Not So Weird: Entanglement of Robust and Non-Robust Features in Neural Network Classifiers. Preprint. 2021.
- **Springer, Jacob M**; Reinstadler, Bryn Marie; O'Reilly, Una-May. STRATA: Simple, Gradient-Free Attacks for Models of Code. 3rd Workshop on Adversarial Learning Methods for Machine Learning and Data Mining @ KDD. 2021.
- **Springer, Jacob M**; Kenyon, Garrett T. It's Hard for Neural Networks To Learn the Game of Life. International Joint Conference on Neural Networks (IJCNN). 2021.
- Wang, Daniel A; Strauss, Charles MS; **Springer, Jacob M**; Thresher, Austin; Pritchard, Howard; Kenyon, Garrett T. Sparse MP4. IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI). 2020.
- **Springer, Jacob M**; Strauss, Charles S; Thresher, Austin M; Kim, Edward; Kenyon, Garrett T. Classifiers based on deep sparse coding architectures are robust to deep learning transferable examples. Preprint. 2018.
- **Springer, Jacob M**; Feng, Wu-chang. Teaching with Angr: A Symbolic Execution Curriculum and CTF. USENIX Workshop on Advances in Security Education (ASE). 2018.

Professional & Research Experience

Carnegie Mellon University

PhD Student

○ Advised by Aditi Raghunathan

Pittsburgh, PA

August 2022 – Present

Cold Spring Harbor Laboratory

Machine Learning & Computational Neuroscience Research

Cold Spring Harbor, NY

January 2022 – July 2022

- Advised by Anthony Zador

Los Alamos National Laboratory

Machine Learning & Computational Neuroscience Research

- Advised by Garrett Kenyon

MIT

Machine Learning Research

- Advised by Una-May O'Reilly

Portland State University

Computer Security Education Research

- Advised by Wu-chang Feng

Los Alamos, NM
June 2018 – December 2021

Cambridge, MA
June 2020 – August 2020

Portland, OR
June 2017 – August 2017

Awards & Achievements

- Hertz Fellowship, Finalist, 2023
- NSF Graduate Research Fellowship, 2022
- Barry M. Goldwater Scholarship, 2020
- Finalist, National Merit Scholarship, 2017