Sam Greydanus

sam.greydanus@gmail.com • greydanus.github.io • Google Scholar • github.com/greydanus

I study how neural networks learn and generalize. My past work focused on physics, interpretability, and AI for science. I have written 25 papers with 4,400+ citations, including *Hamiltonian Neural Networks* and *Lagrangian Neural Networks*. Since 2024 I have built healthcare AI systems now used by national insurers and a federal Medicare program. I am a research collaborator at General Intuition, where I train transformers for world models.

EDUCATION

Dartmouth College, Hanover, NH 2017

BA in Physics, magna cum laude. Physics Chair Prize. GPA 3.81/4.0.

Honors thesis with James Whitfield (high honors): neural networks that find many-body quantum ground states.

CURRENT RESEARCH

Transformers Scale Differently on Tasks with 1D, 2D, and 3D Structure *with Pim de Witte, General Intuition*

Studying how a transformer's error grows with problem size on tasks with 1D, 2D and 3D structure, and which spatial priors stop that growth. In early results, one model shows a substantial improvement over current methods on CLRS-30 and several other medium-scale tasks. Trained on about 40 H100 GPUs. Paper in preparation.

PenGPT: A Generative Transformer for Handwriting and Sketches *independent, 2025–2026*

Small transformers (5M and 54M parameters) that output pen strokes. They write in 1,192 handwriting styles and sketch from text prompts, including scenes that are not in the training data. Paper in preparation.

SELECTED PUBLICATIONS

Hamiltonian Neural Networks *NeurIPS 2019*

S. Greydanus, M. Dzamba, J. Yosinski

Neural networks that learn to conserve energy from data. 2,000+ citations and 80+ named variants in later papers.

Lagrangian Neural Networks *ICLR 2020 workshop (DeepDiffEq)*

M. Cranmer, **S. Greydanus**, S. Hoyer, P. Battaglia, D. Spergel, S. Ho

Learns arbitrary Lagrangians without needing canonical coordinates. 1,000+ citations.

Visualizing and Understanding Atari Agents *ICML 2018 (short oral)*

S. Greydanus, A. Koul, J. Dodge, A. Fern

Perturbation-based saliency maps that explain deep RL agents. 500+ citations; covered by *Science* (2018).

Scaling Down Deep Learning with MNIST-1D *ICML 2024*

S. Greydanus, D. Kobak

A tiny benchmark for the science of deep learning; used in 10 notebooks of Prince's *Understanding Deep Learning*.

Neural Reparameterization Improves Structural Optimization *NeurIPS 2019 workshop (oral)*

S. Hoyer, J. Sohl-Dickstein, **S. Greydanus**

Parameterizing a design with a neural network yields better structures. Basis of US Patent 11,574,093.

EXPERIENCE

General Intuition, *Part-time research collaborator at a world-models lab (see above)* 2026–present

Humata Health, acquired by R1 in 2026 2024–present

Lead AI Scientist; first engineering hire on the payer product

- Built the company's AI system for health-insurance prior authorization through four architecture generations. Automation rose from 0 to 95%; cost per case fell 50x. CMS chose it for Medicare's WISeR model in Oklahoma.
- Built the evaluation pipeline: hundreds of clinician-labeled datasets, regression tests, and fairness audits that validate model, prompt, and guideline changes.

Windscape AI, *Lead ML Scientist; adapted the GraphCast weather model to wind-farm data* 2022–2024

MSS Engineering, *Drafting Engineer* 2021

Oregon State University and ML Collective, *Independent researcher: Lagrangian NNs, MNIST-1D* 2019–2020

Google Brain, San Francisco, CA 2018–2019
Google AI Resident

- First author of Hamiltonian Neural Networks (NeurIPS 2019).
- With Chris Olah: The Paths Perspective on Value Learning (*Distill*); associate editor at *Distill*.
- With Jascha Sohl-Dickstein and Stephan Hoyer: Neural Reparameterization Improves Structural Optimization (NeurIPS workshop oral, later patented) and meta-learned update rules (Cosyne).

Oregon State University, Corvallis, OR 2017–2018
Researcher, DARPA Explainable AI program, with Alan Fern

- Interpreted deep RL agents using saliency maps (ICML 2018) and finite-state machines (ICLR 2019).

Microsoft, *Software engineering intern, Azure Machine Learning* 2016

CERN, *Particle physics intern, Higgs simulations (Mendes group)* 2015

ADDITIONAL PUBLICATIONS

- The Cursive Transformer. **S. Greydanus**, Z. Wimpee. *arXiv, 2025*.
- Nature's Cost Function: Simulating Physics by Minimizing the Action. T. Strang, I. Caruso, **S. Greydanus**. *ICLR 2023 workshop (Physics4ML)*.
- A Tutorial on Structural Optimization. **S. Greydanus**. *arXiv, 2022*.
- Dissipative Hamiltonian Neural Networks. A. Sosanya, **S. Greydanus**. *arXiv, 2022*.
- Piecewise-constant Neural ODEs. **S. Greydanus**, S. Lee, A. Fern. *arXiv, 2021*.
- Self-classifying MNIST Digits. E. Randazzo, A. Mordvintsev, E. Niklasson, M. Levin, **S. Greydanus**. *Distill, 2020*.
- The Story of Airplane Wings. **S. Greydanus**. *arXiv, 2020*.
- The Paths Perspective on Value Learning. **S. Greydanus**, C. Olah. *Distill, 2019*.
- Learning Finite State Representations of Recurrent Policy Networks. A. Koul, **S. Greydanus**, A. Fern. *ICLR 2019*.
- Meta-Learning Biologically Plausible Semi-Supervised Update Rules. K. Gu, **S. Greydanus**, L. Metz, N. Maheswaranathan, J. Sohl-Dickstein. *Cosyne 2019*.
- Learning the Enigma with Recurrent Neural Networks. **S. Greydanus**. *arXiv, 2017*.

OTHER

- **Mentoring (first papers):** A. Sosanya (Dissipative HNNs), T. Strang (Nature's Cost Function), Z. Wimpee (Cursive Transformer). Two went on to their first software engineering jobs, one to his employer's research team.
- **Invited talks:** Princeton University (Lagrangian Neural Networks, Feb 2020); Emory University workshop *Can Machine Learning Learn New Physics?* (June 2020).
- **Open source:** 1,900+ GitHub stars, including hamiltonian-nn (550), scribe (290; handwriting synthesis) and mnist1d (264; 20,000 PyPI downloads in the past six months).
- **Patent:** US 11,574,093 B2 (granted 2023, Google), with S. Hoyer and J. Sohl-Dickstein.
- **Writing:** research blog at greydanus.github.io; two *Distill* articles.
- **Leadership:** co-founded and co-captained the first US team at the International Physicists' Tournament (2017).