

Jascha Sohl-Dickstein, PhD

Google Scholar GitHub *Email:* jascha.sohldickstein@gmail.com

PROFESSIONAL EXPERIENCE

Principal Scientist (2023-present), Senior Staff Research Scientist (2020-2023), Staff Research Scientist (2017-2020), Senior Research Scientist (2015-2017)
Google Brain, Google DeepMind

- Currently co-lead 38 person research team

Visiting Scholar (2012-2015)
Stanford University, Applied Physics

Academic Resident (2012-2014)
Khan Academy (online education nonprofit)

Research Associate, Mars Exploration Rovers (2001-2005)
NASA Jet Propulsion Laboratory

EDUCATION

PhD, Biophysics, May 2012
University of California at Berkeley

BA, Physics, May 2001
BA, College Scholar program (in Quantum Computing), May 2001
Cornell University

SELECTED OPEN SOURCE PROJECTS

Beyond the Imitation Game benchmark (BIG-bench) – github.com/google/BIG-bench

- A large scale, diverse, collaborative benchmark intended to probe large language models, and extrapolate their future capabilities.
- I co-led this project, wrote the original design doc, was the top GitHub contributor

Neural Tangents – github.com/google/neural-tangents

- The only high-level neural network library for defining, training, and evaluating both finite and *infinite* width neural networks.

SELECTED PUBLICATIONS

h-index: 51
i10-index: 89
Citations: 21,000

Diffusion models (including the original paper introducing the model family):

J Sohl-Dickstein, EA Weiss, N Maheswaranathan, S Ganguli. *Deep unsupervised learning using nonequilibrium thermodynamics*. International Conference on Machine Learning (2015).

Y Song, J Sohl-Dickstein, DP Kingma, A Kumar, S Ermon, B Poole. *Score-Based Generative Modeling through Stochastic Differential Equations*. International Conference on Learning Representations (2021). **Oral, Outstanding Paper award (top 0.25%)**

Theory of infinite width neural networks (we showed that infinite width neural networks are linear in their parameters throughout training, and that their output is a Gaussian process):

J Lee, L Xiao, SS Schoenholz, Y Bahri, R Novak, J Sohl-Dickstein, J Pennington. *Wide neural networks of any depth evolve as linear models under gradient descent*. Neural Information Processing Systems (2019).

J Lee, Y Bahri, R Novak, SS Schoenholz, J Pennington, J Sohl-Dickstein. *Deep neural networks as gaussian processes*. International Conference on Learning Representations (2018).

Adversarial examples transfer from machines to humans (!) (motivated by AI safety concerns):

V Veerabadrán, et al. *Subtle adversarial image manipulations influence both human and machine perception*. to appear, Nature Communications (2023).

GF Elsayed, S Shankar, B Cheung, N Papernot, A Kurakin, I Goodfellow, J Sohl-Dickstein. *Adversarial examples that fool both computer vision and time-limited humans*. Neural Information Processing Systems (2018).

Learned optimizers (given enough data and compute, learned functions outperform hand-designed heuristics – this will also hold for learned parameter update functions):

L Metz, J Harrison, CD Freeman, A Merchant, L Beyer, J Bradbury, N Agrawal, B Poole, I Mordatch, A Roberts, J Sohl-Dickstein *VeLO: Training Versatile Learned Optimizers by Scaling Up*. arXiv preprint (2022)

P Vicol, L Metz, J Sohl-Dickstein. *Unbiased Gradient Estimation in Unrolled Computation Graphs with Persistent Evolution Strategies*. International Conference on Machine Learning (2021). **Long oral, Outstanding Paper award (top 0.02%)**

L Metz, N Maheswaranathan, J Nixon, CD Freeman, J Sohl-Dickstein. *Understanding and correcting pathologies in the training of learned optimizers*. International Conference on Machine Learning (2019). **Long oral (top 5%)**

L Metz, N Maheswaranathan, B Cheung, J Sohl-Dickstein. *Meta-Learning Update Rules for Unsupervised Representation Learning*. International Conference on Learning Representations (2019). **Oral (top 1.5%)**

Other influential papers:

L Soderblom, et al. *Pancam multispectral imaging results from the Opportunity rover at Meridiani Planum*. Science (2004).

L Soderblom, et al. *Pancam multispectral imaging results from the Spirit rover at Gusev crater*. Science (2004).

J Sohl-Dickstein, P Battaglino, M DeWeese. *Minimum probability flow learning*. International Conference on Machine Learning (2011). **Distinguished Paper award (top 4 submissions)**

G Tucker, A Mnih, CJ Maddison, J Lawson, J Sohl-Dickstein. *Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models*. Advances in Neural Information Processing Systems (2017). **Oral (top 1.2%)**

L Dinh, J Sohl-Dickstein, S Bengio. *Density estimation using Real NVP*. International Conference on Learning Representations (2017).

L Metz, B Poole, D Pfau, J Sohl-Dickstein. *Unrolled generative adversarial networks*. International Conference on Learning Representations (2017).

C Piech, J Spencer, J Huang, S Ganguli, M Sahami, L Guibas, J Sohl-Dickstein. *Deep knowledge tracing*. Neural Information Processing Systems (2015).

R Burnell, et al. *Rethink reporting of evaluation results in AI*. Science (2023).