

Work Experience

Research Scientist, Blueshift	DeepMind	Mar 2019 – present
<ul style="list-style-type: none"> Reasoning in LLMs (e.g. for math, coding, and the sciences), deep reinforcement learning 		
Research Intern, Neuroscience	DeepMind	Jun 2017 – Oct 2017
<ul style="list-style-type: none"> Multi-task reinforcement learning, variational information bottleneck 		
Machine Learning Intern	Spotify	Jun 2016 – May 2017
<ul style="list-style-type: none"> Probabilistic models of musical taste, Bayesian hypothesis testing 		
Data Science Intern	Zynga	Jun 2015 – Aug 2015
<ul style="list-style-type: none"> Supervised learning on imbalanced datasets 		

Education

PhD, Physics	Princeton University	2012 – 2018
<ul style="list-style-type: none"> <i>Research</i>: information-theoretic regularization in supervised, unsupervised, and reinforcement learning <i>Advisors</i>: David J Schwab, William Bialek <i>Awards</i>: Hertz Fellowship, Department of Energy Computational Sciences Graduate Fellowship 		
MPhil, Information Engineering	University of Cambridge	2011 – 2012
<ul style="list-style-type: none"> <i>Research</i>: neural network models for dendritic integration of synaptic inputs <i>Advisor</i>: Máté Lengyel <i>Awards</i>: Churchill Scholarship 		
BA, Physics and BS, Math	University of Southern California	2006 – 2011
<ul style="list-style-type: none"> <i>Research</i>: quantum algorithms, quantum information theory, computational neuroscience <i>Advisors</i>: Bartlett Mel, Paolo Zanardi, Andrew Childs <i>Awards</i>: USC Presidential Scholarship, Order of the Laurel and the Palm 		

Select Publications¹

-
- Gemini Team. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arxiv*, 2024.
 - Aaditya K. Singh & **DJ Strouse**. Tokenization counts: the impact of tokenization on arithmetic in frontier LLMs. *arxiv*, 2024.
 - Ted Moskovitz, Aaditya Singh, **DJ Strouse**, Tuomas Sandholm, Ruslan Salakhutdinov, Anca D. Dragan, & Stephen McAleer. Confronting reward model overoptimization with constrained RLHF. *International Conference on Learning Representations (ICLR)*, 2024.
 - Michael Laskin, Luyu Wang, Junhyuk Oh, Emilio Parisotto, Stephen Spencer, Richie Steigerwald, **DJ Strouse**, Steven Hansen, Angelos Filos, Ethan Brooks, Maxime Gazeau, Himanshu Sahni, Satinder Singh, & Vlad Mnih. In-context Reinforcement Learning with Algorithm Distillation. *International Conference on Learning Representations (ICLR)*, 2023.
 - DJ Strouse**,* Kate Baumli, David Warde-Farley, Vlad Mnih, & Steven Hansen.* Learning more skills through optimistic exploration. *International Conference on Learning Representations (ICLR)*, 2022.
 - Allison Tam, Neil Rabinowitz, Andrew Lampinen, Nicholas A Roy, Stephanie Chan, **DJ Strouse**, Jane Wang, Andrea Banino, & Felix Hill. Semantic exploration from language abstractions and pretrained representations. *Neural Information Processing Systems (NeurIPS)*, 2022.
 - DJ Strouse**,* Kevin R. McKee, Matt Botvinick, Edward Hughes, & Richard Everett.* Collaborating with Humans without Human Data. *Neural Information Processing Systems (NeurIPS)*, 2021.
 - DJ Strouse** & David Schwab. The information bottleneck and geometric clustering. *Neural Computation (NECO)*, 2019.
 - DJ Strouse**, Max Kleiman-Weiner, Josh Tenenbaum, Matt Botvinick, & David Schwab. Learning to share and hide intentions using information regularization. *Neural Information Processing Systems (NIPS)*, 2018.
 - DJ Strouse** & David Schwab. The deterministic information bottleneck. *Neural Computation (NECO)*, 2017.

¹ See www.djstrouse.com for latest project and publication information.