

CONTACT INFORMATION	353 Serra Mall Stanford, CA 94305	neiswanger@cs.stanford.edu (503) 464-6152
RESEARCH INTERESTS	I develop machine learning methods to perform efficient optimization and experimental design in costly real-world settings, where resources are limited. My work spans topics in active learning, uncertainty quantification, Bayesian decision making, and reinforcement learning. I apply these methods downstream to solve problems in science and engineering, for example in the physical sciences and machine learning systems.	
CURRENT POSITION	<b>Stanford University</b> , Stanford, CA, United States <i>Postdoctoral Scholar, Computer Science Department</i>	<b>September 2020 – Present</b>
	Working with Stefano Ermon, and affiliated with the the StatsML Group, Stanford AI Lab, and SLAC National Accelerator Laboratory.	
EDUCATION	<b>Carnegie Mellon University</b> , Pittsburgh, PA, United States <i>Ph.D. in Machine Learning</i>	<b>September 2012 – August 2019</b>
	Advisor: Eric Xing Thesis: Post-Inference Methods for Scalable Probabilistic Modeling and Sequential Decision Making	
	<b>Columbia University</b> , New York, NY, United States <i>Bachelor of Science</i>	<b>September 2008 – May 2012</b>
	Major: Applied Mathematics, Minor: Computer Science Research Advisors: Chris Wiggins and Frank Wood	
RESEARCH EXPERIENCE	<b>Stanford University</b> , Stanford, CA, United States <i>Probabilistic Methods, Decision Making, Uncertainty Quantification</i>	<b>September 2020 – Present</b>
	Working with Stefano Ermon on probabilistic methods for sequential decision making under uncertainty, experimental design, and uncertainty quantification, applied to science and ML systems.	
	<b>Carnegie Mellon University</b> , Pittsburgh, PA, United States <i>Scalable Probabilistic Modeling and Sequential Decision Making</i>	<b>September 2012 – 2020</b>
	Worked with Eric Xing to develop scalable inference and optimization algorithms, and with Jeff Schneider and Barnabas Poczos to develop methods for sequential decision making under uncertainty.	
	<b>Columbia University</b> , New York, NY, United States <i>Statistics and Computer Vision Research</i>	<b>March 2011 – May 2012</b>
	Worked under the supervision of Frank Wood to develop probabilistic models and inference algorithms for unsupervised detection, tracking, and summarization of objects in videos.	
	<i>Machine Learning and Cell Biology Research</i>	<b>October 2010 – July 2012</b>
	Worked under the supervision of Chris Wiggins to develop software to detect, track, and analyze the motility of T cells for members of the Michael Dustin Laboratory at the NYU Skirball Institute.	
PUBLICATIONS	<b>Conference and Journal Publications</b>	
	<i>*Denotes equal contribution. †Denotes senior authorship.</i>	
	W. Neiswanger*, L. Yu*, S. Zhao, C. Meng, S. Ermon. “Generalizing Bayesian Optimization with Decision-theoretic Entropies.” <i>Conference on Neural Information Processing Systems (NeurIPS)</i> , 2022.	
	V. Mehta, I. Char, J. Abbate, R. Conlin, M. Boyer, S. Ermon, J. Schneider, W. Neiswanger†. “Exploration via Planning for Information about the Optimal Trajectory.” <i>Conference on Neural Information Processing Systems (NeurIPS)</i> , 2022.	

- Y. Xiao, P. Liang, U. Bhatt, W. Neiswanger, R. Salakhutdinov, L. P. Morency. “Uncertainty Quantification with Pre-trained Language Models: A Large-Scale Empirical Analysis.” *Findings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2022.
- J. Song, L. Yu, W. Neiswanger, S. Ermon. “A General Recipe for Likelihood-free Bayesian Optimization.” *International Conference on Machine Learning (ICML)*, 2022.
- C. Marx, S. Zhao, W. Neiswanger, S. Ermon. “Modular Conformal Calibration.” *International Conference on Machine Learning (ICML)*, 2022.
- V. Mehta, B. Paria, J. Schneider, S. Ermon, W. Neiswanger<sup>†</sup>. “An Experimental Design Perspective on Model-Based Reinforcement Learning.” *International Conference on Learning Representations (ICLR)*, 2022.
- M. Fenstermacher et al. “DIII-D research advancing the physics basis for optimizing the tokamak approach to fusion energy.” *Nuclear Fusion*, 2022.
- C. Meng, E. Liu, W. Neiswanger, J. Song, M. Burke, D. Lobell, S. Ermon. “IS-COUNT: Large-scale Object Counting from Satellite Images with Covariate-based Importance Sampling.” *AAAI Conference on Artificial Intelligence (AAAI)*, 2022.
- A. Qiao, S. Choe, S. Subramanya, W. Neiswanger, Q. Ho, H. Zhang, G. Ganger, E. Xing. “Pollux: Co-adaptive Cluster Scheduling for Goodput-Optimized Deep Learning.” *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2021. **Recipient of the Jay Lepreau Best Paper Award.**
- Y. Chung, W. Neiswanger, I. Char, J. Schneider. “Beyond Pinball Loss: Quantile Methods for Calibrated Uncertainty Quantification.” *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- A. Narayan, P. Molino, K. Goel, W. Neiswanger, C. Re. “Personalized Benchmarking with the Ludwig Benchmarking Toolkit.” *Datasets and Benchmarks Track, Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- Y. Liu, S. Khandagale, C. White, W. Neiswanger. “Synthetic Benchmarks for Scientific Research in Explainable Machine Learning.” *Datasets and Benchmarks Track, Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- W. Neiswanger, K. Wang, S. Ermon. “Bayesian Algorithm Execution: Estimating Computable Properties of Black-box Functions Using Mutual Information.” *International Conference on Machine Learning (ICML)*, 2021.
- K. Tran\*, W. Neiswanger\*, K. Broderick, E. Xing, J. Schneider, Z. Ulissi. “Computational catalyst discovery: Active classification through myopic multiscale sampling.” *The Journal of Chemical Physics*, 2021.
- W. Neiswanger, A. Ramdas. “Uncertainty quantification using martingales for misspecified Gaussian processes.” *32nd International Conference on Algorithmic Learning Theory (ALT)*, 2021.
- B. Boecking, W. Neiswanger, E. Xing, A. Dubrawski. “Interactive Weak Supervision: Learning Useful Heuristics for Data Labeling.” *International Conference on Learning Representations (ICLR)*, 2021.
- V. Mehta, I. Char, W. Neiswanger, Y. Chung, A. Nelson, M. Boyer, E. Kolemen, J. Schneider. “Neural Dynamical Systems.” *60th IEEE Conference on Decision and Control (CDC)*, 2021.
- C. White, W. Neiswanger, Y. Savani. “BANANAS: Bayesian Optimization with Neural Architectures for Neural Architecture Search.” *AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- C. White, W. Neiswanger, S. Nolen, Y. Savani. “A Study on Encodings for Neural Architecture Search.” *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- K. Kandasamy, K. Vysyaraju, W. Neiswanger, B. Paria, C. Collins, J. Schneider, B. Póczos, E. Xing. “Tuning Hyperparameters without Grad Students: Scalable and Robust Bayesian Optimisation with Dragonfly.” *Journal of Machine Learning Research (JMLR)*, 2020.

- K. Tran\*, W. Neiswanger\*, J. Yoon, E. Xing, Z. Ulissi. "Methods for Comparing Uncertainty Quantifications for Material Property Predictions." *Machine Learning: Science and Technology (MLST)*, 2020.
- K. Korovina, S. Xu, K. Kandasamy, W. Neiswanger, B. Póczos, J. Schneider, E. Xing. "ChemBO: Bayesian Optimization of Small Organic Molecules with Synthesizable Recommendations." *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- I. Char, Y. Chung, W. Neiswanger, K. Kandasamy, A. Nelson, M. Boyer, E. Kolemen, J. Schneider. "Offline Contextual Bayesian Optimization." *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- K. Kandasamy, W. Neiswanger, R. Zhang, A. Krishnamurthy, J. Schneider, B. Póczos. "Myopic Posterior Sampling for Adaptive Goal Oriented Design of Experiments." *International Conference on Machine Learning (ICML)*, 2019.
- V. Mayya, E. Judokusumo, E. Abu-Shah, W. Neiswanger, C. Sachar, D. Depoil, L. Kam and M. L. Dustin. "Cutting Edge: Synapse Propensity of Human Memory CD8 T Cells Confers Competitive Advantage over Naive Counterparts." *The Journal of Immunology*, 2019.
- K. Kandasamy, W. Neiswanger, J. Schneider, B. Póczos, E. Xing. "Neural Architecture Search with Bayesian Optimisation and Optimal Transport." *Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
- V. Mayya, E. Judokusumo, E. Abu Shah, C. Peel, W. Neiswanger, D. Depoil, D. Blair, C. Wiggins, L. Kam, and M. Dustin. "Durable interactions of T cells with T cell receptor stimuli in the absence of a stable immunological synapse." *Cell Reports*, 2018.
- W. Neiswanger, E. Xing. "Post-Inference Prior Swapping." *International Conference on Machine Learning (ICML)*, 2017.
- R. Steorts, M. Barnes, W. Neiswanger. "Performance Bounds for Graphical Record Linkage." *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.
- F. Caron\*, W. Neiswanger\*, F. Wood, A. Doucet, M. Davy. "Generalized Pólya Urn for Time-Varying Pitman-Yor Processes." *Journal of Machine Learning Research (JMLR)*, 2017.
- Y. Wang, V. Sadhanala, W. Dai, W. Neiswanger, S. Sra, E. Xing. "Parallel and Distributed Block-Coordinate Frank-Wolfe Algorithms." *International Conference on Machine Learning (ICML)*, 2016.
- J. Oliva, W. Neiswanger, B. Póczos, J. Schneider, E. Xing. "Fast Function to Function Regression." *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- W. Neiswanger, C. Wang, E. Xing. "Asymptotically Exact, Embarrassingly Parallel MCMC." *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2014.
- W. Neiswanger, C. Wang, Q. Ho, E. Xing. "Modeling Citation Networks using Latent Random Offsets." *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2014.
- W. Neiswanger, F. Wood, E. Xing. "The Dependent Dirichlet Process Mixture of Objects for Detection-free Tracking and Object Modeling." *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2014.
- J. Oliva, W. Neiswanger, B. Póczos, J. Schneider, E. Xing. "Fast Distribution to Real Regression." *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2014.
- V. Mayya\*, W. Neiswanger\*, R. Medina, C. Wiggins, and M. Dustin. "Integrative analysis of T cell motility from multi-channel microscopy data using TIAM." *Journal of Immunological Methods*, 2014.
- V. Mayya, W. Neiswanger, D. Blair, C. Wiggins, M. Dustin. "Characterization of synapse-kinapse balance in CD8 T cells." *The Journal of Immunology*, 2012.

## Preprints

S. Miskovich\*, W. Neiswanger\*, W. Colocho, C. Emma, J. Garrahan, T. Maxwell, C. Mayes, S. Ermon, A. Edelen, D. Ratner. “Bayesian Algorithm Execution for Tuning Particle Accelerator Emittance with Partial Measurements.” *In Submission*, 2022.

S. Choe, W. Neiswanger, P. Xie, E. Xing. “Betty: An Automatic Differentiation Library for Multilevel Optimization.” *In Submission*, 2022.

B. Boecking, N. Roberts, W. Neiswanger, S. Ermon, F. Sala, A. Dubrawski. “Generative Modeling Helps Weak Supervision (and Vice Versa).” *In Submission*, 2022.

X. Li, V. Mehta, J. Kirschner, I. Char, W. Neiswanger, J. Schneider, A. Krause, I. Bogunovic. “Near-optimal Policy Identification in Active Reinforcement Learning.” *In Submission*, 2022.

## Technical Reports and Workshop Publications

S. Choe, W. Neiswanger, P. Xie, E. Xing. “Betty: An Automatic Differentiation Library for Multilevel Optimization.” *NeurIPS Workshop on Meta-Learning*, 2022.

R. Tu, N. Roberts, V. Prasad, S. Nayak, P. Jain, F. Sala, G. Ramakrishnan, A. Talwalkar, W. Neiswanger, C. White. “AutoML for Climate Change: A Call to Action.” *NeurIPS Workshop on Tackling Climate Change with Machine Learning*, 2022.

V. Mehta, I. Char, J. Abbate, R. Conlin, M. Boyer, S. Ermon, J. Schneider, W. Neiswanger<sup>†</sup>. “Sample-efficient Plasma Control by Planning for Optimal Trajectory Information.” *ICML Workshop on Adaptive Experimental Design and Active Learning in the Real World*, 2022.

Y. Chung, I. Char, H. Guo, J. Schneider, W. Neiswanger<sup>†</sup>. “Uncertainty toolbox: an open-source library for assessing, visualizing, and improving uncertainty quantification.” *ICML Workshop on Uncertainty and Robustness in Deep Learning*, 2021.

B. Lengerich, W. Neiswanger, E. Lengerich, E. Xing. “Disentangling Increased Testing From Covid-19 Epidemic Spread.” *medRxiv 2020.07.09.20141762*, 2020.

Y. Chung, I. Char, W. Neiswanger, K. Kandasamy, A. Nelson, M. Boyer, E. Kolemen, J. Schneider. “Offline Contextual Bayesian Optimization for Nuclear Fusion.” *NeurIPS Workshop on Machine Learning and the Physical Sciences*, 2019.

K. Korovina, S. Xu, K. Kandasamy, W. Neiswanger, B. Póczos, J. Schneider, E. Xing. “ChemBO: Bayesian Optimization of Small Organic Molecules with Synthesizable Recommendations.” *NeurIPS Workshop on Machine Learning and the Physical Sciences*, 2019.

C. White, W. Neiswanger, Y. Savani. “Deep Uncertainty Estimation for Model-based Neural Architecture Search.” *NeurIPS Workshop on Bayesian Deep Learning*, 2019.

C. White, W. Neiswanger, Y. Savani. “Neural Architecture Search via Bayesian Optimization with a Neural Network Prior.” *NeurIPS Workshop on Meta-Learning*, 2019.

W. Neiswanger, X. Liu, E. Xing, “Low Communication Distributed Black Box VI”, *The International Conference on Probabilistic Programming (PROBPROG)*, 2018.

K. Kandasamy, W. Neiswanger, R. Zhang, A. Krishnamurthy, J. Schneider, B. Póczos. “Sequential Bayesian Design of Experiments via Probabilistic Programming”, *The International Conference on Probabilistic Programming (PROBPROG)*, 2018.

W. Neiswanger, E. Xing, “Prior Swapping for Data-Independent Inference”, *ICML Workshop on Data-Efficient Machine Learning*, 2016.

W. Neiswanger, C. Wang, E. Xing, “Embarrassingly Parallel Variational Inference.” *NeurIPS Workshop on Advances in Variational Inference*, 2014.

Y. Wang, V. Sadhanala, W. Dai, W. Neiswanger, S. Sra, E. Xing, “Asynchronous Parallel Block-Coordinate Frank-Wolfe”, *NeurIPS Workshop on Optimization for Machine Learning*, 2014.

W. Neiswanger, C. Wang, E. Xing, “Embarrassingly Parallel MCMC via Density Product Estimation”, *NIPS Workshop on Randomized Methods for Machine Learning*, 2013.

W. Neiswanger and F. Wood, “Unsupervised Detection and Tracking of Multiple Objects with Dependent Dirichlet Process Mixtures”, *New York Academy of Sciences, Machine Learning Symposium*, New York, NY. October 21, 2011.

#### AWARDS

Jay Lepreau Best Paper Award, OSDI, 2021

ICML Reviewer Award, 2020

AWS Machine Learning Research Award, 2019

Pittsburgh Filmmakers Spring Mix, Audience Choice Award, 2017

ICML Reviewer Award, 2015

Columbia University Department of Applied Mathematics Faculty Award, 2012

KAUST International Research Competition, 1st Place, 2012

CUSP Summer Enhancement Fellowship Winner 2009, 2011

Columbia C.P. Davis Scholar, 2008

Intel ISEF Grand Prize Award Winner, 4th Place 2006, 3rd Place 2007, 4th Place 2008

#### INVITED TALKS

*AI-Guided Experiments with Bayesian Algorithm Execution*. Toyota Research Institute (TRI), October 2022.

*AI-Guided Experiments for ML Systems, AutoML, and Science*. 1st CASL Project Workshop, October 2022.

*Bayesian Algorithm Execution for Tuning Particle Accelerator Emittance with Partial Measurements*. General Assembly of the Helmholtz International Laboratory (HIR<sup>3</sup>X), September 2022.

*BAX, InfoBAX, and Applications to Experimental Design and Reinforcement Learning*. Oxford University Deep-Prob Seminar, February 2022.

*Going Beyond Global Optima with Bayesian Algorithm Execution*. AutoML Seminars, European Lab for Learning and Intelligent Systems (ELLIS), December 2021.

*AI Meets Biology Panel*. State-Of-The-Art Conference, December 2021.

*InfoBAX: Estimating Black-box Functions Using Mutual Information*. London Data Science Journal Club, November 2021.

*Going Beyond Global Optima with Bayesian Algorithm Execution*. Saint Petersburg State University, Industrial Mathematics Seminar, September 2021.

*Uncertainty quantification in machine learning with applications to sequential decision making*. Korea University, Machine learning and Policy Studies International Workshop, September 2021.

*Going Beyond Global Optima with Bayesian Algorithm Execution*. University College London (UCL), Statistical Machine Learning Group, July 2021.

*Going Beyond Global Optima with Bayesian Algorithm Execution*. LBNL Camera Workshop on Autonomous Discovery in Science and Engineering, April 2021.

*Going Beyond Global Optima with Bayesian Algorithm Execution*. AI Seminar, SLAC National Accelerator Laboratory, April 2021.

*Working with Scientists to Develop Uncertainty Models for Improved Sequential Decision Making.* SIAM CSE21 Minisymposium, March 2021.

*Probabilistic Programming for Sequential Decision Making Under Uncertainty.* Stanford University, March 2020.

*Machine Learning Bandits.* Carnegie Mellon University TechNights Program, March 2018.

*Post-Inference Prior Swapping.* Machine Learning Lunch Seminar. Carnegie Mellon University, November 2017.

*Variational Bayesian Inference and Mean Field Approximations.* Guest lecture, Probabilistic Graphical Models Class. Carnegie Mellon University, March 2017.

*Embarrassingly Parallel MCMC.* Machine Learning Lunch Seminar. Carnegie Mellon University, October 2014.

*Asymptotically Exact, Embarrassingly Parallel MCMC.* Oxford Computational Statistics / Statistical Machine Learning Reading Group. Oxford University, July 2014.

*Asymptotically Exact, Embarrassingly Parallel MCMC.* Department of Information Engineering Tea Talk. Oxford University, July 2014.

*Unsupervised Detection and Tracking of Multiple Objects with Dependent Dirichlet Process Mixtures.* Tutorial for Statistics Department Data Mining Class. Columbia University, January 2012.

TEACHING

*Guest Lecture for Probabilistic Graphical Models Class at CMU* **Spring 2017**  
Gave a guest lecture on variational inference and mean field approximations.

*Two Guest Lectures for Probabilistic Graphical Models Class at CMU* **Spring 2015**  
Gave two lectures on approximate inference algorithms and probabilistic latent variable models.

*Teaching Assistant for CMU 10-601: Introduction to Machine Learning* **Spring 2015**  
Led recitation lectures, mentored project groups, designed tests and homework assignments.

*Teaching Assistant for CMU 10-708: Probabilistic Graphical Models* **Spring 2014**  
Led recitation lectures, mentored project groups, designed tests and homework assignments.

*Guest Lecture for Data Mining and Machine Learning Class at Columbia University* **Spring 2012**  
Gave a guest lecture on nonparametric Bayesian models for computer vision applications.

ACTIVITIES AND  
PROFESSIONAL  
SERVICE

*Workshop Organizer* **July 2022**  
Co-organized the ICML 2022 Workshop on Adaptive Experimental Design and Active Learning in the Real World.

*Workshop Organizer* **July 2021**  
Co-organized the ICML 2021 Workshop on Machine Learning for Data.

*CASL Project* **September 2020 – Present**  
Co-founded and organized the CASL Project for composable, automatic, and scalable ML systems.

*Workshop Organizer* **July 2020**  
Co-organized the ICML 2020 Workshop on Real World Experiment Design and Active Learning.

*Reviewer* **June 2013 – Present**  
NeurIPS, ICML, JMLR, ICLR, AISTATS, UAI, AAAI, IJCAI, CVPR, TKDE, JSIG, AISM, Statistics and Computing, Digital Signal Processing, PLOS ONE, JASA, Biometrika, Journal of the Royal Statistical Society: Series B.

*Workshop Organizer*

**December 2019**

Co-organized the NeurIPS 2019 Workshop on Learning with Rich Experience (LIRE).

*CMU AI+ Club*

**January 2018 – August 2020**

Co-founded Carnegie Mellon's AI+ club. Helped organize and host events.

*CMU Machine Learning Lunch Seminar*

**September 2013 – August 2019**

Organized Carnegie Mellon's Machine Learning Lunch weekly seminar.

*Columbia University Beginner Machine Learning Reading Group*

**June 2011 – May 2012**

Founder, participant, and presenter. Helped organize and conduct meetings.

*Society for Industrial and Applied Mathematics (SIAM)*

**May 2011 – May 2012**

Member of SIAM and board member of the Columbia University SIAM student chapter.

*Tutoring*

**June 2010 – May 2012**

Tutored high school students in mathematics (calculus, precalculus, trigonometry).

PROGRAMMING

Python, R, SQL, Shell Scripting,  $\LaTeX 2_{\epsilon}$ .

SOFTWARE AND  
SYSTEMS

*Uncertainty Toolbox*

<https://uncertainty-toolbox.github.io>

A Python toolbox for predictive uncertainty quantification, calibration, metrics, and visualization.

*Betty*

<https://github.com/leopard-ai/betty>

An automatic differentiation library for generalized meta-learning and multilevel optimization.

*AdaptDL*

<https://github.com/petuum/adaptdl>

Resource-adaptive cluster scheduler for deep learning training.

*TorchUQ*

<https://github.com/TorchUQ/torchuq>

A library for uncertainty quantification based on PyTorch.

*Naszilla*

<https://github.com/naszilla/naszilla>

A Python library for neural architecture search (NAS).

*Dragonfly*

<https://github.com/dragonfly/dragonfly>

An open source Python library for scalable Bayesian optimisation.