

CONTACT INFORMATION	6105 Gates Center, 5000 Forbes Ave. Pittsburgh, PA 15213	willie@cs.cmu.edu (503) 464 6152
RESEARCH INTERESTS	I develop methods for distributed inference and probabilistic programming to help scale and automate machine learning. My interests include computer vision, computational biology, and analysis of text and network data. I also work on sequential decision making under uncertainty, with application to problems in the sciences and engineering.	
EDUCATION	Carnegie Mellon University , Pittsburgh, PA, United States <i>Ph.D. in Machine Learning</i> August 2012 – Present <ul style="list-style-type: none">• Advisor: Eric Xing• Thesis: <i>Post-Inference Methods for Scalable Probabilistic Modeling</i> Columbia University , New York, NY, United States <i>Bachelor of Science</i> September 2008 – May 2012 <ul style="list-style-type: none">• Major: Applied Mathematics, Minor: Computer Science• Research Advisors: Chris Wiggins and Frank Wood	
RESEARCH EXPERIENCE	Carnegie Mellon University , Pittsburgh, PA, United States <i>Scalable Machine Learning for Probabilistic Models</i> September 2012 – Present Worked under the supervision of Eric Xing to develop scalable inference and optimization algorithms, and models, for tasks in computer vision, natural language processing, and network analysis. Columbia University , New York, NY, United States <i>Statistics and Machine Vision Research</i> March 2011 – May 2012 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> October 2010 – July 2012 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. <i>Statistics and Political Science Research</i> May 2009 – January 2010 Worked with members of Statistics and Political Science departments to study prison sentence lengths.	
WORK EXPERIENCE	Petuum , Pittsburgh, PA, United States <i>Research Scientist Intern</i> Oct 2017 – Present Developed inference algorithms and probabilistic programming systems for distributed Bayesian learning, which can be deployed at scale and in an automated fashion on diverse computing platforms. Bitly , New York, NY, United States <i>Data Science Intern</i> May 2012 – August 2012 Worked on methods to predict real-world events with bitly data, and on network models to analyze diffusion of information in the social web. Gained experience with Hadoop for large data analysis. Memorial Sloan Kettering Cancer Center , New York, NY, United States <i>Development Analytics Intern</i> May 2010 – August 2010 Carried out analysis, modeling, and forecasting of donor behavior in order to optimize donor targeting. Gained experience with SAS, Enterprise Miner, SQL, R, and Tableau.	

AWARDS

Pittsburgh Filmmakers Spring Mix, Audience Choice Award, 2017
ICML Reviewer Award, 2015
KAUST International Research Poster Competition for Undergraduates, 1st Place, 2012
Department of Applied Mathematics Faculty Award, 2012
CUSP Summer Enhancement Fellowship Winner 2009, 2011
Columbia C.P. Davis Scholar, 2008
Intel International Science and Engineering Fair (ISEF) Finalist, 2006, 2007, 2008
Intel ISEF Grand Prize Award Winner, 4th Place 2006, 3rd Place 2007, 4th Place 2008
National Merit Scholar, 2008

PUBLICATIONS

Conference and Journal Publications

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

W. Neiswanger, F. Caron, 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, A. Dubey, C. Wang, E. Xing, “Embarrassingly Parallel Variational Inference in Nonconjugate Models”, *arXiv:1510.04163*, 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

W. Neiswanger, V. Mayya, 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”, *Journal of Immunology*, id: 169548, vol: 188, 2012

Preprints

K. Kandasamy, W. Neiswanger, J. Schneider, B. Póczos, E. Xing. “Neural Architecture Search with Bayesian Optimisation and Optimal Transport”.

K. Kandasamy, W. Neiswanger, R. Zhang, A. Krishnamurthy, J. Schneider, B. Póczos. “Myopic Bayesian Design of Experiments via Posterior Sampling and Probabilistic Programming”.

Workshop Publications

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”, *NIPS 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”, *NIPS 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

V. Mayya, W. Neiswanger, C. Wiggins, and M. Dustin, “Characterization of naïve CD8 T cell motility during antigen signaling”, *19th Annual International Cancer Immunotherapy Symposium*, New York, NY. October 3-5, 2011

W. Neiswanger and F. Wood, “Time Dependent Dirichlet Process Mixture Models for Multiple Target Tracking”, *Machine Learning Summer School Poster Session*, West Lafayette, IN. June 17, 2011

M. Dewar, W. Neiswanger, V. Mayya, M. Dustin, and Chris Wiggins, “Tracking T-Cell Motility”, *Nanomedicine Center for Mechanobiology, NIH Nanomedicine Development Centers 5th Annual Awardee Meeting*, Bethesda, MD. March 28-29, 2011

INVITED TALKS

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

ACTIVITIES AND PROFESSIONAL SERVICE

Reviewer **June 2013 – Present**
Reviewed publications for NIPS, ICML, JMLR, AISTATS, UAI, AAAI, IJCAI, TKDE, STCO, DSP, JSIG, AISM, PLOS ONE, JASA, Biometrika.

CMU Machine Learning Lunch Seminar **September 2013 – Present**
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).

RELEVANT CLASSES Machine Learning, Data Mining, Graph Theory and Combinatorics, Applied Computational Math and Physics, Numerical Methods, Scientific Computation, Probability, Statistics, Real Analysis, Complex Analysis, Abstract Algebra, Data Driven Modeling, Statistical Machine Learning, Probabilistic Graphical Models, Randomized Algorithms, Optimization, Advanced Optimization and Randomized Methods

PROGRAMMING Python, Matlab, Java, R, SQL, SAS, Mathematica, C, C++, Shell Scripting, L^AT_EX 2_ε.