LESTER MACKEY

Microsoft Research New England   1 Memorial Dr., Cambridge, MA 02474
lmackey@microsoft.com   https://web.stanford.edu/~lmackey

EDUCATION

University of California, Berkeley, Ph.D., May 2012
Computer Science, GPA: 4.0/4.0
Designated Emphasis in Communication, Computation, and Statistics

University of California, Berkeley, M.A., December 2011
Statistics, GPA: 4.0/4.0

Princeton University, B.S.E., summa cum laude, June 2007
Computer Science, GPA: 3.98/4.0
Certificate in Applied and Computational Mathematics

EXPERIENCE

Principal Researcher, Microsoft Research New England, 9/2019 – Present
Adjunct Professor of Statistics, Stanford University, 9/2016 – Present
Researcher, Microsoft Research New England, 9/2016 – 8/2019
Assistant Professor of Statistics and (by courtesy) of Computer Science, Stanford University, 9/2013 – 9/2016
Simons Math+X Postdoctoral Fellow, Stanford University, 9/2012 – 8/2013
Research Intern, Google Inc., Summer 2011
Graduate Student Researcher, University of California, Berkeley, Summer 2009, 2010
Recommender System Architect, Umamibud (later Discovereads, acquired by Goodreads, acquired by Amazon), Summer 2008
Research Intern, AT&T Labs, Summer 2007
Research Intern, Princeton University, Summer 2006
Software Design Engineer Intern, Microsoft Small Business Accounting team, Summer 2005
Research Intern, Intel Strategic CAD Labs, Integrated Design and Verification team, Summer 2004

AWARDS & HONORS

MacArthur Fellowship, 2023
Ethel Newbold Prize (Awarded biennially to an outstanding statistical scientist for excellence in research in mathematical statistics), Bernoulli Society, 2023
Elected Fellow, Institute of Mathematical Statistics, 2023
Outstanding Paper Award, NeurIPS for "Gradient Estimation with Discrete Stein Operators," 2022
Elected to the COPSS Leadership Academy, Committee of Presidents of Statistical Societies, 2022
Student Paper Award, ASA SCSG for "Distribution Compression in Near-linear Time" with Abhishek Shetty and Raaz Dwivedi, 2022
Winner of U.S. Bureau of Reclamation's $525K Subseasonal Climate Forecast Rodeo (first place for 2-4 week temperature forecasting, second place for 4-6 week precipitation forecasting, first place for temperature hindcasting), 2019
Featured in Computer History Museum's exhibit on "Dreamers, Builders, and Thinkers," 2019
Frederick E. Terman Fellowship, 2013
San Francisco Business Times 40 under 40 Emerging Leaders, 2013
First Place in the $50K ALS Prediction Prize4Life Challenge for predicting Lou Gehrig's disease progression, 2012
Eli Jury Award (Outstanding contribution to Systems, Communications, and Control), U.C. Berkeley, 2012
Outstanding Graduate Student Instructor Award, U.C. Berkeley, 2011
Best Student Paper Award, ICML for "On the Consistency of Ranking Algorithms," 2010
Second Place in the $1 million Netflix Prize competition for collaborative filtering, 2009
National Science Foundation Graduate Research Fellowship, 2009
National Defense Science and Engineering Graduate Fellowship, 2009
SIGPLAN CACM Research Highlights Nomination for "Fault-tolerant Typed Assembly Language," 2008
Best Paper Award, ACM SIGPLAN PLDI for "Fault-tolerant Typed Assembly Language," 2007
James Hayes-Edgar Palmer Prize in Engineering, Princeton U., 2007
Philip Y. Goldman '86 Senior Prize in Computer Science, Princeton U., 2007
Sigma Xi Scientific Research Society, 2007

U.C. Berkeley Chancellor's Fellowship, 2007
AT&T Labs Fellowship, 2007
Tau Beta Pi Engineering Honor Society, 2007
M. Taylor Pyne Honor Prize (Highest general distinction awarded to an undergraduate at Princeton U.), 2007
Computing Research Association Outstanding Undergraduate Award Winner, 2007
Accenture Award in Computer Science, 2006
Phi Beta Kappa Honor Society, 2006
Barry M. Goldwater Scholarship, 2006
Shapiro Prize for Academic Excellence, Princeton U., 2004, 2005
Microsoft Scholar, 2004, 2005, 2006
Top 20 of ACM International Collegiate Programming Contest Regional Semifinals, 2005
Princeton U. Freshman First Honor Prize (Awarded to one student for exceptional academic achievement in freshman year), 2004
Quin Morton '36 Writing Seminar Essay Prize (One of ten best essays produced by freshmen at Princeton U.), 2004
Creativity Foundation Legacy Award Winner, 2003-2004
Intel Science Talent Search National Winner, 2003
Intel Science Talent Search Glenn T. Seaborg Award Winner, 2003
Intel International Science and Engineering Fair Operation Cherry Blossom Award, 2003
Intel International Science and Engineering Fair 2nd Place Grand Award in Mathematics, 2003
Valedictorian, Half Hollow Hills High School West, 2003

WORK IN SUBMISSION

Debiased Distribution Compression. Lingxiao Li, Raaz Dwivedi, and Lester Mackey.

SatCLIP: Global, General-Purpose Location Embeddings with Satellite Imagery. Konstantin Klemmer, Esther Rolf, Caleb Robinson, Lester Mackey, and Marc Rußwurm.

Self-Taught Optimizer (STOP): Recursively Self-Improving Code Generation. Eric Zelikman, Eliana Lorch, Lester Mackey, and Adam Tauman Kalai.

Reflections from the Workshop on AI-Assisted Decision Making for Conservation. Lily Xu, Esther Rolf, Sara Beery, Joseph R. Bennett, Tanya Berger-Wolf, Tanya Birch, Elizabeth Bondi-Kelly, Justin Brashares, Melissa Chapman, Anthony Corso, Andrew Davies, Nikhil Garg, Angela Gaylard, Robert Heilmayr, Hannah Kerner, Konstantin Klemmer, Vipin Kumar, Lester Mackey, Claire Monteleoni, Paul Moorcroft, Jonathan Palmer, Andrew Perrault, David Thau, and Milind Tambe.

Controlling Moments with Kernel Stein Discrepancies. Heishiro Kanagawa, Arthur Gretton, and Lester Mackey.

Budget-Constrained Bounds for Mini-Batch Estimation of Optimal Transport. David Alvarez-Melis, Nicolo Fusi, Lester Mackey, and Tal Wagner.

Efficient Concentration with Gaussian Approximation. Morgane Austern and Lester Mackey.

Near-optimal inference in adaptive linear regression. Koulik Khamaru, Yash Deshpande, Lester Mackey, and Martin J. Wainwright.

Weighted Meta-Learning. Diana Cai, Rishit Sheth, Lester Mackey, and Nicolo Fusi.

Teacher-Student Compression with Generative Adversarial Networks. Ruishan Liu, Nicolo Fusi, and Lester Mackey.

PEER-REVIEWED PUBLICATIONS

Kernel Thinning. Raaz Dwivedi and Lester Mackey. *Journal of Machine Learning Research (JMLR).* To appear.

Targeted Separation and Convergence with Kernel Discrepancies. Alessandro Barp, Carl-Johann Simon-Gabriel, Mark Girolami, and Lester Mackey. *Journal of Machine Learning Research (JMLR).* To appear.

Do Language Models Know When They're Hallucinating References? Ayush Agrawal, Lester Mackey, and Adam Tauman Kalai. *Findings of the Association for Computational Linguistics: EACL.* Mar. 2024.

SubseasonalClimateUSA: A Dataset for Subseasonal Forecasting and Benchmarking. Soukayna Mouatadid, Paulo Orenstein, Genevieve Flaspohler, Miruna Oprescu, Judah Cohen, Franklyn Wang, Sean Knight, Maria Geogdzhayeva, Sam Levang, Ernest Fraenkel, and Lester Mackey. *Advances in Neural Information Processing Systems (NeurIPS).* Dec. 2023.

Should I Stop or Should I Go: Early Stopping with Heterogeneous Populations. Hammaad Adam, Fan Yin, Mary Hu, Neil Tenenholtz, Lorin Crawford, Lester Mackey, and Allison Koenecke. *Advances in Neural Information Processing Systems (NeurIPS).* Dec. 2023.

Learning Rate Free Bayesian Inference in Constrained Domains. Louis Sharrock, Lester Mackey, and Christopher Nemeth. *Advances in Neural Information Processing Systems (NeurIPS).* Dec. 2023.

A Finite-Particle Convergence Rate for Stein Variational Gradient Descent. Jiaxin Shi and Lester Mackey. *Advances in Neural Information Processing Systems (NeurIPS).* Dec. 2023.

Bounding Wasserstein distance with couplings. Niloy Biswas and Lester Mackey. *Journal of the American Statistical Association (JASA).* Nov. 2023.

Evaluation of Dependency Structure for Multivariate Weather Predictors using Copulas. Samuel C. Maina, Dorcas Mwigereri, Jonathan Weyn, Lester Mackey, and Millicent Ochieng. ACM Journal on Computing and Sustainable Societies (COMPASS). Sep. 2023.

Metrizing Weak Convergence with Maximum Mean Discrepancies. Carl-Johann Simon-Gabriel, Alessandro Barp, and Lester Mackey. *Journal of Machine Learning Research (JMLR).* August 2023.

Adaptive Bias Correction for Improved Subseasonal Forecasting. Soukayna Mouatadid, Paulo Orenstein, Genevieve Flaspohler, Judah Cohen, Miruna Oprescu, Ernest Fraenkel, and Lester Mackey. *Nature Communications.* June 2023.

A Kernel Stein Test for Comparing Latent Variable Models. Heishiro Kanagawa, Wittawat Jitkrittum, Lester Mackey, Kenji Fukumizu, and Arthur Gretton. *Journal of the Royal Statistical Society, Series B (JRSSB).* May 2023.

Independent finite approximations for Bayesian nonparametric inference. Tin D. Nguyen, Jonathan H. Huggins, Lorenzo Masoero, Lester Mackey, and Tamara Broderick. *Bayesian Analysis.* May 2023.

Compress Then Test: Powerful Kernel Testing in Near-linear Time. Carles Domingo-Enrich, Raaz Dwivedi, and Lester Mackey. *International Conference on Artificial Intelligence and Statistics (AISTATS).* April 2023.

Stein's Method Meets Statistics: A Review of Some Recent Developments. Andreas Anastasiou, Alessandro Barp, François-Xavier Briol, Bruno Ebner, Robert E. Gaunt, Fatemeh Ghaderinezhad, Jackson Gorham, Arthur Gretton, Christophe Ley, Qiang Liu, Lester Mackey, Chris. J. Oates, Gesine Reinert, and Yvik Swan. *Statistical Science.* February 2023.

Social Norm Bias: Residual Harms of Fairness-Aware Algorithms. Myra Cheng, Maria De-Arteaga, Lester Mackey, and Adam Tauman Kalai. *WIREs Data Mining and Knowledge Discovery.* Jan. 2023.

Gradient Estimation with Discrete Stein Operators. Jiaxin Shi, Yuhao Zhou, Jessica Hwang, Michalis K. Titsias, and Lester Mackey. *Advances in Neural Information Processing Systems (NeurIPS).* Dec. 2022. NeurIPS 2022 Outstanding Paper Award.

Optimal Thinning of MCMC Output. Marina Riabiz, Wilson Chen, Jon Cockayne, Pawel Swietach, Steven A. Niederer, Lester Mackey, and Chris. J. Oates. *Journal of the Royal Statistical Society, Series B (JRSSB).* Sep. 2022.

Scalable Spike-and-Slab. Niloy Biswas, Lester Mackey, and Xiao-Li Meng. *International Conference on Machine Learning (ICML).* July 2022.

Distribution Compression in Near-linear Time. Abhishek Shetty, Raaz Dwivedi, and Lester Mackey. *International Conference on Learning Representations (ICLR).* April 2022. 2022 ASA SCSG Student Paper Award.

Generalized Kernel Thinning. Raaz Dwivedi and Lester Mackey. *International Conference on Learning Representations (ICLR).* April 2022.

Sampling with Mirrored Stein Operators. Jiaxin Shi, Chang Liu, and Lester Mackey. *International Conference on Learning Representations (ICLR)*. April 2022.

An Open Repository of Real-Time COVID-19 Indicators. Alex Reinhart, Logan Brooks, Maria Jahja, Aaron Rumack, Jingjing Tang, Wael Al Saeed, Taylor Arnold, Amartya Basu, Jacob Bien, Ángel A. Cabrera, Andrew Chin, Eu Jing Chua, Brian Clark, Nat DeFries, Jodi Forlizzi, Samuel Gratzl, Alden Green, George Haff, Robin Han, Addison J. Hu, Sangwon Hyun, Ananya Joshi, Jimi Kim, Andrew Kuznetsov, Wichada La Motte-Kerr, Yeon Jin Lee, Kenneth Lee, Zachary C. Lipton, Michael X. Liu, Lester Mackey, Kathryn Mazaitis, Daniel J. McDonald, Balasubramanian Narasimhan, Natalia L. Oliveira, Pratik Patil, Adam Perer, Collin A. Politsch, Samyak Rajanala, Dawn Rucker, Nigam H. Shah, Vishnu Shankar, James Sharpnack, Dmitry Shemetov, Noah Simon, Vishakha Srivastava, Shuyi Tan, Robert Tibshirani, Elena Tuzhilina, Ana Karina Van Nortwick, Valérie Ventura, Larry Wasserman, Jeremy C. Weiss, Kristin Williams, Roni Rosenfeld, and Ryan J. Tibshirani. *Proceedings of the National Academy of Sciences (PNAS)*. December 2021.

DeepMiner: Discovering Interpretable Representations for Mammogram Classification and Explanation. Jimmy Wu, Bolei Zhou, Diondra Peck, Scott Hsieh, Vandana Dialani, Lester Mackey, and Genevieve Patterson. *Harvard Data Science Review (HDSR)*. October 2021.

Online Learning with Optimism and Delay. Genevieve Flaspohler, Francesco Orabona, Judah Cohen, Soukayna Mouatadid, Miruna Oprescu, Paulo Orenstein, and Lester Mackey. *International Conference on Machine Learning (ICML)*. July 2021.

Knowledge Distillation as Semiparametric Inference. Tri Dao, Govinda M. Kamath, Vasilis Syrgkanis, and Lester Mackey. *International Conference on Learning Representations (ICLR)*. May 2021.

Initialization and Regularization of Factorized Neural Layers. Mikhail Khodak, Neil Tenenholtz, Lester Mackey, and Nicolo Fusi. *International Conference on Learning Representations (ICLR)*. May 2021.

Cross-validation Confidence Intervals for Test Error. Pierre Bayle, Alexandre Bayle, Lucas Janson, and Lester Mackey. *Advances in Neural Information Processing Systems (NeurIPS)*. December 2020.

Stochastic Stein Discrepancies. Jackson Gorham, Anant Raj, and Lester Mackey. *Advances in Neural Information Processing Systems (NeurIPS)*. December 2020.

Minimax Estimation of Conditional Moment Models. Nishanth Dikkala, Greg Lewis, Lester Mackey, and Vasilis Syrgkanis. *Advances in Neural Information Processing Systems (NeurIPS)*. December 2020.

Single Point Transductive Prediction. Nilesh Tripuraneni and Lester Mackey. *International Conference on Machine Learning (ICML)*. July 2020.

Approximate Cross-validation: Guarantees for Model Assessment and Selection. Ashia Wilson, Maximilian Kasy, and Lester Mackey. *International Conference on Artificial Intelligence and Statistics (AISTATS)*. June 2020.

Importance Sampling via Local Sensitivity. Anant Raj, Cameron Musco, and Lester Mackey. *International Conference on Artificial Intelligence and Statistics (AISTATS)*. June 2020.

Evaluation of Combined Artificial Intelligence and Radiologist Assessment to Interpret Screening Mammograms. Thomas Schaffter, Diana S. M. Buist, Christoph I. Lee, Yaroslav Nikulin, Dezső Ribli, Yuanfang Guan, William Lotter, Zequn Jie, Hao Du, Sijia Wang, Jiashi Feng, Mengling Feng, Hyo-Eun Kim, Francisco Albiol, Alberto Albiol, Stephen Morrell, Zbigniew Wojna, Mehmet Eren Ahsen, Umar Asif, Antonio Jimeno Yepes, Shivanthan Yohanandan, Simona Rabinovici-Cohen, Darvin Yi, Bruce Hoff, Thomas Yu, Elias Chaibub Neto, Daniel L. Rubin, Peter Lindholm, Laurie R. Margolies, Russell Bailey McBride, Joseph H. Rothstein, Weiva Sieh, Rami Ben-Ari, Stefan Harrer, Andrew Trister, Stephen Friend, Thea Norman, Berkman Sahiner, Fredrik Strand, Justin Guinney, Gustavo Stolovitzky, and the DM DREAM Consortium (including Lester Mackey). *Journal of the American Medical Association Network Open*. March 2020.

Minimum Stein Discrepancy Estimators. Alessandro Barp, Francois-Xavier Briol, Andrew B. Duncan, Mark Girolami, and Lester Mackey. *Advances in Neural Information Processing Systems (NeurIPS)*. December 2019.

Stochastic Runge-Kutta Accelerates Langevin Monte Carlo and Beyond. Xuechen Li, Denny Wu, Lester Mackey, and Murat A. Erdogdu. *Advances in Neural Information Processing Systems (NeurIPS)*. December 2019.

LESTER MACKEY

Accelerating Rescaled Gradient Descent: Fast Optimization of Smooth Functions. Ashia Wilson, Lester Mackey, and Andre Wibisono. *Advances in Neural Information Processing Systems (NeurIPS).* December 2019.

Measuring Sample Quality with Diffusions. Jackson Gorham, Andrew B. Duncan, Sebastian J. Vollmer, and Lester Mackey. *Annals of Applied Probability, 29(5), 2884-2928.* October 2019.

Improving Subseasonal Forecasting in the Western U.S. with Machine Learning. Jessica Hwang, Paulo Orenstein, Judah Cohen, Karl Pfeiffer, and Lester Mackey. *International Conference on Knowledge Discovery and Data Mining (KDD).* August 2019.

Stein Point Markov Chain Monte Carlo. Wilson Ye Chen, Alessandro Barp, Francois-Xavier Briol, Jackson Gorham, Mark Girolami, Lester Mackey, and Chris. J. Oates. *International Conference on Machine Learning (ICML).* July 2019.

Stratification of amyotrophic lateral sclerosis patients: a crowdsourcing approach. Robert Kueffner, Neta Zach, Maya Bronfeld, Raquel Norel, Nazem Atassi, Venkat Balagurusamy, Barbara DiCamillo, Adriano Chio, Merit Cudkowicz, Donna Dillenberger, Javier Garcia-Garcia, Orla Hardiman, Bruce Hoff, Joshua Knight, Melanie L. Leitner, Guang Li, Lara Mangravite, Thea Norman, Liuxia Wang, The ALS Stratification Consortium (including Lester Mackey), Jinfeng Xiao, Wen-Chieh Fang, Jian Peng, Chen Yang, Huan-Jui Chang, and Gustavo Stolovitzky. *Scientific Reports.* January 2019.

A Multifactorial Model of T Cell Expansion and Durable Clinical Benefit in Response to a PD-L1 Inhibitor. Mark DM Leiserson, Vasilis Syrgkanis, Amy Gilson, Miroslav Dudik, Dean F Bajorin, Jonathan Rosenberg, Samuel Funt, Alexandra Snyder, and Lester Mackey. *PLOS One.* December 2018.

S2S reboot: An argument for greater inclusion of machine learning in subseasonal to seasonal forecasts. Judah Cohen, Dim Coumou, Jessica Hwang, Lester Mackey, Paulo Orenstein, Sonja Totz, and Eli Tziperman. *WIREs Climate Change.* December 2018.

Global Non-convex Optimization with Discretized Diffusions. Murat A. Erdogdu, Lester Mackey, and Ohad Shamir. *Advances in Neural Information Processing Systems (NeurIPS).* December 2018.

Random Feature Stein Discrepancies. Jonathan H. Huggins and Lester Mackey. *Advances in Neural Information Processing Systems (NeurIPS).* December 2018.

Stein Points. Wilson Ye Chen, Lester Mackey, Jackson Gorham, Francois-Xavier Briol, and Chris J. Oates. *International Conference on Machine Learning (ICML).* July 2018.

Orthogonal Machine Learning: Power and Limitations. Lester Mackey, Vasilis Syrgkanis, and Ilias Zadik. *International Conference on Machine Learning (ICML).* July 2018.

Accurate Inference for Adaptive Linear Models. Yash Deshpande, Lester Mackey, Vasilis Syrgkanis, and Matt Taddy. *International Conference on Machine Learning (ICML).* July 2018.

Measuring Sample Quality with Kernels. Jackson Gorham and Lester Mackey. *International Conference on Machine Learning (ICML).* August 2017.

Improving Gibbs Sampler Scan Quality with DoGS. Ioannis Mitliagkas and Lester Mackey. *International Conference on Machine Learning (ICML).* August 2017.

Empirical Bayesian Analysis of Simultaneous Changepoints in Multiple Data Sequences. Zhou Fan and Lester Mackey. *Annals of Applied Statistics.* December 2017.

Predicting Patient Cost Blooms in Denmark: A Longitudinal Population-based Study. Suzanne Tamang, Lars Pedersen, Henrik Toft Sørensen, Arnold Milstein, Lester Mackey, Jean-Raymond Betterton, Lucas Janson, and Nigam Shah. *BMJ Open.* January 2017.

Boosted jet tagging with jet-images and deep neural networks. Michael Kagan, Luke de Oliveira, Lester Mackey, Benjamin Nachman, and Ariel Schwartzman. *EPJ Web of Conferences.* November 2016.

LESTER MACKEY

Image Processing, Computer Vision, and Deep Learning: new approaches to the analysis and physics interpretation of LHC events. Ariel Schwartzman, Michael Kagan, Lester Mackey, Benjamin Nachman, and Luke de Oliveira. *Journal of Physics: Conference Series.* October 2016.

Predicting inpatient clinical order patterns with probabilistic topic models vs. conventional order sets. Jonathan H. Chen, Mary K. Goldstein, Steven M. Asch, Lester Mackey, and Russ B. Altman. *Journal of the American Medical Informatics Association*. September 2016.

Efron-Stein Inequalities for Random Matrices. Daniel Paulin, Lester Mackey, and Joel A. Tropp. *Annals of Probability*. September 2016.

Multivariate Stein Factors for a Class of Strongly Log-concave Distributions. Lester Mackey and Jackson Gorham. *Electronic Communications in Probability.* September 2016.

Jet-Images -- Deep Learning Edition. Luke de Oliveira, Michael Kagan, Lester Mackey, Benjamin Nachman, and Ariel Schwartzman. *Journal of High Energy Physics*. July 2016.

Fuzzy Jets. Lester Mackey, Benjamin Nachman, Ariel Schwartzman, and Conrad Stansbury. *Journal of High Energy Physics*. June 2016.

Measuring Sample Quality with Stein's Method. Jackson Gorham and Lester Mackey. *Advances in Neural Information Processing Systems (NeurIPS)*. December 2015.

Weighted Classification Cascades for Optimizing Discovery Significance in the HiggsML Challenge. Lester Mackey, Jordan Bryan, and Man Yue Mo. *Proceedings of the NeurIPS Workshop on High Energy Physics, Machine Learning, and the HiggsML Data Challenge.* August 2015.

Distributed Matrix Completion and Robust Factorization. Lester Mackey, Ameet Talwalkar, and Michael I. Jordan. *Journal of Machine Learning Research*. April 2015.

Crowdsourced analysis of clinical trial data to predict amyotrophic lateral sclerosis progression. Robert Küffner, Neta Zach, Raquel Nore, Johann Hawe, David Schoenfeld, Liuxia Wang, Guang Li, Lilly Fang, Lester Mackey, Orla Hardiman, Merit Cudkowicz, Alexander Sherman, Gokhan Ertaylan, Moritz Grosse-Wentrup, Torsten Hothorn, Jules van Ligtenberg, Jakob H. Macke, Timm Meyer, Bernhard Schölkopf, Linh Tran, Rubio Vaughan, Gustavo Stolovitzky, and Melanie L. Leitner. *Nature Biotechnology*. November 2014.

Combinatorial Clustering and the Beta Negative Binomial Process. Tamara Broderick, Lester Mackey, John Paisley, and Michael I. Jordan. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. April 2014.

Matrix Concentration Inequalities via the Method of Exchangeable Pairs. Lester Mackey, Michael I. Jordan, Richard Y. Chen, Brendan Farrell, and Joel A. Tropp. *Annals of Probability*, 42(3), 906-945. March 2014.

Corrupted Sensing: Novel Guarantees for Separating Structured Signals. Rina Foygel and Lester Mackey. *IEEE Transactions on Information Theory*, 60(2), 1223-1247. February 2014.

Distributed Low-rank Subspace Segmentation. Ameet Talwalkar, Lester Mackey, Yadong Mu, Shih-Fu Chang, and Michael I. Jordan. *IEEE International Conference on Computer Vision (ICCV)*. December 2013.

The Asymptotics of Ranking Algorithms. John C. Duchi, Lester Mackey, and Michael I. Jordan. *Annals of Statistics*, 41(5), 2292-2323. November 2013.

Joint Link Prediction and Attribute Inference using a Social-Attribute Network. Neil Zhenqiang Gong, Ameet Talwalkar, Lester Mackey, Ling Huang, Eui Chul Richard Shin, Emil Stefanov, Elaine (Runting) Shi, and Dawn Song. *ACM Transactions on Intelligent Systems and Technology (ACM TIST)*. March 2013.

Divide-and-Conquer Matrix Factorization. Lester Mackey, Ameet Talwalkar, and Michael I. Jordan. *Advances in Neural Information Processing Systems (NeurIPS)*. December 2011.

Visually Relating Gene Expression and in vivo DNA Binding Data. Min-Yu Huang, Lester Mackey, Soile Keranen, Gunther Weber, Michael Jordan, David Knowles, Mark Biggin, and Bernd Hamann. *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. November 2011.

Mixed Membership Matrix Factorization. Lester Mackey, David Weiss, and Michael I. Jordan. *International Conference on Machine Learning (ICML)*. June 2010.

On the Consistency of Ranking Algorithms. John Duchi, Lester Mackey, and Michael I. Jordan. *International Conference on Machine Learning (ICML)*. June 2010. ICML 2010 Best Student Paper Award.

Deflation Methods for Sparse PCA. Lester Mackey. *Advances in Neural Information Processing Systems (NeurIPS)*. December 2008.

Fault-tolerant Typed Assembly Language. Frances Perry, Lester Mackey, George A. Reis, Jay Ligatti, David I. August, and David Walker. *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. June 2007. Joint winner of the PLDI 2007 Best Paper Award.

Static Typing for a Faulty Lambda Calculus. David Walker, Lester Mackey, Jay Ligatti, George Reis, and David August. *ACM SIGPLAN International Conference on Functional Programming (ICFP)*. September 2006.

Participatory Design with Proxies: Developing a Desktop-PDA System to Support People with Aphasia. Jordan Boyd-Graber, Sonya Nikolova, Karyn Moffatt, Kenrick Kin, Joshua Lee, Lester Mackey, Marilyn Tremaine, and Maria Klawe. *SIGCHI Conference on Human Factors in Computing Systems (CHI)*. April 2006.

UNPUBLISHED REPORTS

Deriving Matrix Concentration Inequalities from Kernel Couplings. Daniel Paulin, Lester Mackey, and Joel A. Tropp. May 2013.

Feature-Weighted Linear Stacking. Joe Sill, Gabor Takacs, Lester Mackey, and David Lin. November 2009.

INVITED TALKS

Doing Some Good with Machine Learning
- AI for Social Impact Seminar, Harvard University, Apr. 2024.
- AI & ML Community Seminar, Microsoft, Feb. 2024.
- Data Science for Social Change Seminar, Cornell University, Feb. 2023.
- Computer Science Tea, Carleton College, Feb. 2023.
- Keynote - Data Science Week, Purdue University Fort Wayne, Nov. 2022.
- Vector Institute Machine Learning Black and Indigenous Program, Nov. 2022.
- Mathematical Sciences Colloquium, Appalachian State University, Sep. 2022.
- Research Science Institute, Distinguished Guest Lecture, July 2022.
- CIFAR Deep Learning+Reinforcement Learning Summer School, July 2021.
- Statistics Seminar, Luxembourg Statistical Society, June 2021.
- AI Seminar, University of Alberta, May 2021.
- Neyman Seminar, U.C. Berkeley, Dec. 2020.
- Data Science for Social Good Workshop, Georgia Tech, Nov. 2020.
- Keynote - International Conference on Machine Learning (ICML), July 2020.

Advances in Distribution Compression / Kernel Thinning and Stein Thinning
- Stochastics and Statistics Seminar, MIT, December 2023.
- IMSI Workshop on Permutations and Causal Inference, Chicago, August 2023.
- COPSS Leadership Academy Seminar, Joint Statistical Meetings, Toronto, August 2023.
- Ethel Newbold Prize Seminar, World Statistics Congress, Ottawa, July 2023.
- Semi-plenary - Journées de Statistique, French Statistical Society, Brussels, July 2023.
- GRAMSIA: Graphical Models, Statistical Inference, and Algorithms Workshop, Harvard University, May 2023.
- Probability and Statistics Seminar, Boston University, Mar. 2023.

- Information Systems Laboratory Colloquium, Stanford University, Mar. 2023.
- Statistics Seminar, UW Madison, Dec. 2022.
- Data Science Seminar, London School of Economics, May 2022.
- Symposium on Advances in Approximate Bayesian Inference, February 2022.
- Statistics Seminar, University of Michigan, December 2021.
- Statistics and Data Science Seminar, Yale University, November 2021.
- ML Advances and Applications Seminar, Vector Institute & Fields Institute, November 2021.
- ML Theory Seminar, Harvard University, November 2021.
- Computational Bayesian Statistics Journal Club, Flatiron Institute, November 2021.
- Applied Math Colloquium, Cornell University, October 2021.
- Bayesian Data Analysis Webinar, October 2021.
- Statistics Seminar, Cambridge University, October 2021.
- Statistics Colloquium, Columbia University, Sep. 2021.

Adaptive Bias Correction for Improved Subseasonal Forecasting
- Workshop on AI-Assisted Decision-Making for Conservation, Harvard CRCS, Oct. 2022.

Probabilistic Inference and Learning with Stein's Method
- Plenary - SIAM Conference on Mathematics of Data Science (MDS), Sep. 2022.
- Tutorial - Stein's Method: The Golden Anniversary, IMS, Jun. 2022.
- Beyond iid Learning Seminar, ETH Zurich, Dec. 2021.
- Your Model Is Wrong Workshop, NeurIPS, Dec. 2021.
- Stein's Method and Machine Learning Workshop, Royal Statistical Society, Dec. 2021.
- Plenary - International Conference on Monte Carlo Methods and Applications (MCM), Aug. 2021.
- ISI World Statistics Conference, July 2021.
- Probability for Machine Learning Seminar, University of Oxford, Dec. 2020.
- Econometrics and Statistics Colloquium Workshop, U. Chicago, Nov. 2020.
- Joint Statistics and Machine Learning Seminar, Carnegie Mellon University, Nov. 2020.
- Statistics Colloquium, Penn State, Sep. 2020.
- Oberwolfach Uncertainty Quantification Workshop, Oberwolfach, Germany, Mar. 2019.
- Keynote - Data, Learning, and Inference Workshop (DALI), George, South Africa, Jan. 2019.
- Charles Stein Memorial Session, Joint Statistical Meetings, Vancouver, Canada, July 2018.

Cross-validation Confidence Intervals for Test Error
- Inference and Interpretability in a Model-Free Setting, Joint Statistical Meetings, Aug. 2021.

Online Learning with Optimism and Delay
- ICML Time Series Workshop, July 2021.

Knowledge Distillation as Semiparametric Inference
- Directions in ML Speaker Series, Microsoft Research, Apr. 2021.

Improving Subseasonal Forecasting in the Western U.S. with Machine Learning
- ICLR Climate Change AI Workshop, April 2020.
- Keynote - Open Data Science Conference, April 2020.
- NeurIPS Workshop on Tackling Climate Change with Machine Learning, Dec. 2019.
- Statistics and Data Science Conference, MIT, April 2019.
- Computer Science Seminar, Cornell University, November 2018.

Orthogonal Machine Learning: Power and Limitations
- Robust and High-Dimensional Statistics Workshop, Simons Institute for the Theory of Computing, October 2018.

Measuring Sample Quality with Kernels
- Bayes, Machine Learning, and Deep Learning Invited Session, International Society for Bayesian Analysis (ISBA) World Meeting, June 2018.
- Harvard / MIT Econometrics Workshop, MIT, Mar. 2018.

- SAMSI Workshop on Trends and Advances in Monte Carlo Sampling Algorithms, Duke University, Dec. 2017.
- SAMSI Workshop on Quasi-Monte Carlo and High-Dimensional Sampling Methods, Duke University, Aug. 2017.
- Borchard Colloquium on Concentration Inequalities, High Dimensional Statistics, and Stein's Method, Missilac, France, July 2017.
- New England Machine Learning Day, Cambridge, MA, May 2017.
- Machine Learning Seminar, MIT, Mar. 2017.

Measuring Sample Quality with Stein's Method
- Gatsby Unit Seminar, University College London, Oct. 2016.
- Seminar, University of Liege, Sep. 2016.
- Quetelet Seminar, Ghent University, Sep. 2016.
- International Conference on Monte Carlo and Quasi-Monte Carlo Methods, Stanford, CA, Aug. 2016.
- Statistics Seminar, Columbia University, Feb. 2016.
- Quasi-Monte Carlo Invited Session, IMS-ISBA Joint Meeting (MCMSki V), Jan. 2016.
- Wharton Statistics Seminar, University of Pennsylvania, Dec. 2015.
- Neyman Seminar, UC Berkeley, Sep. 2015.
- IMS-Microsoft Research Workshop: Foundations of Data Science, Cambridge, MA, June 2015.
- Stochastics and Statistics Seminar, MIT, May 2015.
- Statistics Seminar, Stanford University, May 2015.

Statistics for Social Good
- AI Now Symposium on the Social and Economic Impact of Artificial Intelligence Technologies, MIT, July 2017.
- Data Science @ Stanford Seminar, Stanford, June 2016.

Matrix Completion and Matrix Concentration
- IDSS Special Seminar, Massachusetts Institute of Technology, Feb. 2016.
- Statistics Seminar, Harvard University, Nov. 2014.
- Blackwell-Tapia Conference, Los Angeles, CA, Nov. 2014.
- Information Systems Laboratory Colloquium, Stanford University, April 2013.
- Statistics Seminar, Yale University, April 2013.
- Statistics Seminar, Columbia University, April 2013.
- Computer Science Seminar, University of Southern California, May 2012.
- Statistics Seminar, Stanford University, Jan. 2012.

Divide-and-Conquer Matrix Factorization
- CS Department Colloquium, Princeton University, Dec. 2015.
- Workshop on Big Data: Theoretical and Practical Challenges, Paris, France, May 2013.
- Kaggle, San Francisco, CA, Feb. 2013.
- Statistical Science Seminar Series, Duke University, Jan. 2012.
- CMS Seminar, Caltech, Jan. 2012.
- SF Bay Area Machine Learning Meetup, San Francisco, CA, Nov. 2011.

Predicting ALS Progression with Bayesian Additive Regression Trees
- Big Data in Biomedicine Conference, Stanford University, May 2015.
- Guest Lecture, Stats 202, Stanford University, Nov. 2013.
- Statistics Seminar, Stanford University, April 2013.
- RECOMB Conference on Regulatory and Systems Genomics, Redwood City, CA, Nov. 2012.

Weighted Classification Cascades for Optimizing Discovery Significance
- NeurIPS Workshop on High-energy particle physics, machine learning, and the HiggsML data challenge (HEPML), December 2014.

Ranking, Aggregation, and You
- Statistics Seminar, University of Chicago, Oct. 2014.
- MacMillan-CSAP Workshop on Quantitative Research Methods, Yale University, Sep. 2014.
- Wharton Statistics Seminar, University of Pennsylvania, Sep. 2014.

- Statistics Seminar, Carnegie Mellon University, Sep. 2014.
- Western Section Meeting, American Mathematical Society, Nov. 2013.
- Statistics Seminar, Stanford University, Sep. 2013.
- Stanford Statistics-Machine Learning Reading Group, Stanford University, Nov. 2012.

Build a Better Netflix, Win a Million Dollars?
- SPARC Camp, Berkeley, CA, Aug. 2014.
- USA Science and Engineering Festival, Washington, DC, Apr. 2012.

Dividing, Conquering, and Mixing Matrix Factorizations
- Technicolor, Palo Alto, CA, June 2013.

Stein's Method for Matrix Concentration
- Institut National de Recherche en Informatique et en Automatique (INRIA), Dec. 2012.
- Berkeley Probability Seminar, University of California, Berkeley, Jan. 2012.

The Story of the Netflix Prize: An Ensembler's Tale
- National Academies' Seminar, Washington, DC, Nov. 2011.

Mixed Membership Matrix Factorization
- Joint Statistical Meetings, Miami Beach, FL, July 2011.

False Event Identification and Beyond: A Machine Learning Approach
- Comprehensive Test Ban Treaty Organization Technical Meeting on Data Mining, Vienna, Austria, Nov. 2009.

The Dinosaur Planet Approach to the Netflix Prize
- LIDS Seminar series, Massachusetts Institute of Technology, Nov. 2008.
- Guest Lecture, Stat 157, University of California, Berkeley, Sept. 2008.
- Process Driven Trading Group, Morgan Stanley, April 2008.

PRESENTATIONS

A Multifactorial Model of T Cell Expansion and Durable Clinical Benefit in Response to a PD-L1 Inhibitor. American Association for Cancer Research (AACR) Annual Meeting, Chicago, IL, Apr. 2018.

Post hoc analysis using PRO-ACT database to evaluate Repository Corticotropin Injection (H.P. Acthar® Gel) as potential treatment for ALS. American Academy of Neurology 70th Annual Meeting, Los Angeles, CA, Apr. 2018.

Divide and Conquer Subspace Segmentation. Neural Information Processing Systems Workshop on Big Learning, Lake Tahoe, NV, Dec. 2012.

Improved Automated Seismic Event Extraction Using Machine Learning. American Geophysical Union Fall Meeting, San Francisco, CA, Dec. 2009.

Machine Learning for Improved Automated Seismic Event Extraction. International Scientific Studies Conference, Vienna, Austria, June 2009.

TEACHING EXPERIENCE

Introduction to Statistical Learning, Stanford University, Winter 2015, 2016
Data Mining and Analysis, Stanford University, Fall 2015
Theory of Statistics, Stanford University, Fall 2013, 2014, 2015
Methods for Applied Statistics: Unsupervised Learning, Stanford University, Spring 2014
Self-paced Center, U.C. Berkeley, Fall 2011
Statistical Learning Theory, U.C. Berkeley, Fall 2010
Practical Machine Learning, U.C. Berkeley, Fall 2009

SERVICE

General Chair, Neural Information and Processing Systems (NeurIPS), 2024
Proceedings of the National Academy of Sciences (PNAS) Statistical and Methodological Review Committee, 2024-2026
COPSS Emerging Leader Award Committee, 2023-2024
IMS Committee on Nominations, 2023-2024
Ethics Review Chair, Neural Information and Processing Systems (NeurIPS), 2023
Census Scientific Advisory Committee, United States Census Bureau, 2021-2022
Diversity and Inclusion Chair, Neural Information and Processing Systems (NeurIPS), 2020, 2021
Science Advisory Board, Institute for Pure & Applied Mathematics (IPAM), 2021
Treasurer, International Society of Bayesian Analysis (ISBA) Section on Bayesian Computation (BayesComp), 2019-2020
Grand Juror, July – Oct. 2019
Organizer, Microsoft Research Colloquium, 2017-2018
Organizer, Stanford Statistics for Social Good, 2013-2016
Organizer, Stanford Statistics Seminar, 2014, 2016
PhD Admissions Committee, Stanford Statistics Department, 2012-2013, 2013-2014, 2015-2016
Stanford Statistics Data Science Committee, 2014-15
Organizer, Stanford Statistics Summer Undergraduate Research Program, 2014, 2015
Stanford Data Science Initiative Working Group, 2014
Computer Science Graduate Student Association Activities Committee, U.C. Berkeley, 2010, 2011
PhD Admissions Committee, U.C. Berkeley Computer Science Division, 2010, 2011

Organizer, Invited Session on Statistical Machine Learning, IMS Annual Meeting, 2022
Organizer, Stein's Method – The Golden Anniversary, Institute of Mathematical Sciences, 2022
Organizer, Advances in Stein's Method and Its Applications in Machine Learning and Optimization, BIRS, 2022
Organizer, Workshop on Stein's Method in Machine Learning and Statistics, ICML, 2019
Organizer, Workshop on AI for Social Good, ICML, 2019
Organizer, Session on Probabilistic Methods in Machine Learning, SPA, 2018
Chair, Session on Convex Modeling for High-Dimensional Data Analysis, ASC-IMS, 2014
Organizer, Workshop on Sparse Representation and Low-rank Approximation, NeurIPS, 2011

Senior Meta-Reviewer for International Conference on Machine Learning (ICML), 2022-
- Area Chair / Meta-Reviewer for ICML, 2019-2021
- Reviewer for ICML, 2013-2018
Senior Area Chair for Neural Information and Processing Systems (NeurIPS), 2021-2023
- Area Chair for NeurIPS, 2017-2020
- Reviewer for NeurIPS, 2010-2016
Editor for Environmental Data Science, 2021-2023
Associate Editor for SIAM Mathematics of Data Science (SIMODS), 2023-
Associate Editor for Annals of Statistics, 2019-2022

Reviewer for Annals of Statistics
Reviewer for Conference on Artificial Intelligence (AAAI)
Reviewer for International Conference on Artificial Intelligence and Statistics (AISTATS)
Reviewer for Bernoulli
Reviewer for Biometrika
Reviewer for Conference on Learning Theory (COLT)
Reviewer for Electronic Journal of Statistics
Reviewer for Foundations and Trends in Machine Learning (FnTML)
Reviewer for IEEE Transactions on Information Theory
Reviewer for International Joint Conferences on Artificial Intelligence (IJCAI)
Reviewer for Journal of Machine Learning Research (JMLR)
Reviewer for Machine Learning Journal
Reviewer for NeurIPS Workshop on Computational Social Science and the Wisdom of Crowds
Reviewer for Neurocomputing
Reviewer for Richard Tapia Celebration of Diversity in Computing Conference, 2021-

PROFESSIONAL AFFILIATIONS

LESTER MACKEY

American Statistical Association (ASA)
Association for the Advancement of Artificial Intelligence (AAAI)
Association of Computing Machinery (ACM)
Bernoulli Society
Institute of Electrical and Electronics Engineers (IEEE)
Institute of Mathematical Statistics (IMS)
International Society for Bayesian Analysis (ISBA)
Society for Industrial and Applied Mathematics (SIAM)