

MATTHEW DOUGLAS HOFFMAN

Senior Research Scientist
Google Research

matt@matthewdhoffman.com
<http://www.matthewdhoffman.com>

EDUCATION:

Ph.D., Computer Science, Princeton University, November 2010.

M.A., Computer Science, Princeton University, September 2006.

B.A., Management-Engineering, Claremont McKenna College, December 2003.

B.S., Computer Science, School of Engineering and Applied Science, Columbia University, May 2003.

WORK EXPERIENCE:

December 2016–Present: **Senior Research Scientist**, Google Research.

November 2014–December 2016: **Senior Research Scientist**, Creative Technologies Laboratory, Adobe Research.

July 2012–October 2014: **Research Scientist**, Creative Technologies Laboratory, Adobe Research.

Conducted research in machine learning and statistics, with applications in audio, vision, digital imaging, social networking, information retrieval, and analytics. Wrote papers, supervised several Ph.D. interns per year, consulted with product teams, transferred technology to product teams (including Photoshop, Behance.net, Adobe Stock, Ink, Character Animator, and Data Workbench), filed patents.

September 2010–June 2012: **Postdoctoral Researcher**, Columbia University Department of Statistics.

Postdoctoral research fellowship under the supervision of Prof. Andrew Gelman, addressing computational issues associated with hierarchical probabilistic models.

August 2004–September 2010: **Research Assistant**, Princeton University Computer Science Department.

Conducted research and wrote on online variational inference, nonparametric Bayesian models for latent source discovery and separation, automatic music annotation, content-based estimation of similarity for recorded music, automatic estimation of dissonance in recorded music, using global optimization heuristics and perceptual modeling to control complex musical synthesizers, content-based search of large speech audio corpora, and improving the auditory display of medical data obtained in real time.

May–August 2008: **Intern**, Google.

Sped up core Google web search retrieval by 3% by further improving highly optimized C++ code.

February–August 2004: **Programmer**, Bear, Stearns and Company, Inc.

Developed interfaces, analytics, and tools in C++ for mortgage trading desk.

Summer 2003: **Intern**, Crushing Music.

Assisted with day-to-day operations of advertising music production house.

Summers 1998–2002: **Intern** at Convergys, Global InfoTek Inc., Decisive Analytics Corporation.

Developed interfaces and implemented algorithms for various Java applications.

AWARDS:

Best poster presentation, New York Academy of Sciences Machine Learning Symposium 2016 (for

“A Variational Analysis of Stochastic Gradient Algorithms”).

Best student paper award, ISMIR 2013 (for “Beta Process Sparse Nonnegative Matrix Factorization for Music”).

2nd place best student paper award, New York Academy of Sciences Machine Learning Symposium 2010 (for “Online Learning for Latent Dirichlet Allocation”).

Best student paper award, ISMIR 2009 (for “Easy as CBA: A Simple Probabilistic Model for Tagging Music”).

2nd place best student paper award, New York Academy of Sciences Machine Learning Symposium 2009 (for “Finding Latent Sources in Recorded Music With a Shift-Invariant HDP”).

TEACHING EXPERIENCE:

Summer 2013: Visiting Lecturer, Stanford University. Statistics 300: Advanced Topics in Statistics.
Spring 2009: Adjunct Professor, New York University. E85.2607: Advanced Digital Signal Theory.
Spring 2007: Teaching Assistant, Princeton University. Computer Science/Music 325: Transforming Reality By Computer (Digital Signal Processing for Music).
Spring 2006: Teaching Assistant, Princeton University. Computer Science 226: Algorithms and Data Structures.
Fall 2005: Teaching Assistant, Princeton University. Computer Science 109: Computers in Our World.

PROFESSIONAL ACTIVITIES:

Area Chair, NIPS 2013, ICML 2015, ICML 2017

Co-organizer, NIPS 2015 workshop on Advances in Approximate Bayesian Inference

Co-organizer, NIPS 2014 workshop on Advances in Variational Inference

Journal Reviewing: Journal of Machine Learning Research (JMLR); Journal of the American Statistical Society (JASA); IEEE Transactions on Audio, Speech, and Language Processing (IEEE TASLP); IEEE Transactions on Neural Networks (IEEE TNN); Pattern Recognition Letters; IEEE Transactions on Pattern Analysis and Machine Intelligence (IEEE TPAMI); Statistics and Computing

Conference Reviewing: Advances in Neural Information Processing Systems (NIPS); International Conference on Machine Learning (ICML); International Conference on Artificial Intelligence and Statistics (AISTATS); Conference on Artificial Intelligence (AAAI); International Joint Conferences on Artificial Intelligence (IJCAI); IEEE Conference on Audio, Speech, and Signal Processing (ICASSP); International Society for Music Information Retrieval Conference (ISMIR); International Computer Music Conference (ICMC).

PUBLICATIONS (see <http://tinyurl.com/mdh-gscholar> for a more up-to-date list and pdf links):

- A. Saedi, M. Hoffman, R. Adams, M. Johnson, “The Segmented iHMM: A Simple, Efficient Hierarchical HMM,” in *Proceedings of the International Conference on Machine Learning*, New York, 2016.
- S. Mandt, M. Hoffman, D. Blei, “A Variational Analysis of Stochastic Gradient Algorithms,” in *Proceedings of the International Conference on Machine Learning*, New York, 2016.
- Z. Liu, Y. Wang, M. Dontcheva, M. Hoffman, S. Walker, A. Wilson, “Patterns and Sequences: Interactive Exploration of Clickstreams to Understand User Paths,” in *IEEE Transactions on Visualization and Computer Graphics*, 2016.
- M. Rudolph, M. Hoffman, A. Hertzmann, “A Joint Model for Who-to-Follow and What-to-View Recommendations on Behance,” in *Proceedings of the International Workshop on Modeling Social Media*, Montreal, 2016.
- J. Regier, A. Miller, J. McAuliffe, R. Adams, M. Hoffman, D. Lang, D. Schlegel, Prabhat, “Celeste: Variational Inference for a Generative Model of Astronomical Images,” in *Proceedings of the International Conference on Machine Learning*, Lille, 2015.
- L. Theis, M. Hoffman, “A Trust-Region Method for Stochastic Variational Inference with Applications to Streaming Data,” in *Proceedings of the International Conference on Machine Learning*, Lille, 2015.
- D. Liang, M. Hoffman, G. Mysore, “Speech Dereverberation Using a Learned Speech Model,” in *Proceedings of the IEEE International Conference on Audio, Speech, and Signal Processing (ICASSP)*, Brisbane, 2015.
- B. Carpenter, A. Gelman, M. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, A. Riddell, “Stan: A Probabilistic Programming Language,” in the *Journal of Statistical Software*, 2015.
- H. Izadinia, B. Russell, A. Farhadi, M. Hoffman, A. Hertzmann, “Deep Classifiers from Image Tags in the Wild,” in *Proceedings of Multimedia COMMONS, ACM Multimedia*, Brisbane, 2015.
- M. Hoffman, D. Blei, “Structured Stochastic Variational Inference,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, San Diego, 2015.
- F. Agostinelli, M. Hoffman, P. Sadowski, P. Baldi, “Learning Activation Functions to Improve Deep Neural Networks,” in *Proceedings of the International Conference on Learning Representations (Workshop track)*, San Diego, 2015.
- M. Hoffman, “A Problem with (and fix for) Variational Bayesian NMF,” in *Proceedings of the IEEE GlobalSIP Conference*, Atlanta, 2014.
- M. Hoffman, A. Gelman, “The No-U-Turn Sampler,” in *Journal of Machine Learning Research*, 2014.
- P. Smaragdis, C. Fevotte, G. Mysore, N. Mohammadiha, M. Hoffman, “Static and Dynamic Source Separation Using Nonnegative Factorizations: A unified view,” in *IEEE Signal Processing Magazine*, 2014.
- D. Liang, M. Hoffman, G. Mysore, D.P.W. Ellis, “Speech Decoloration Based on the Product-of-Filters Model,” in *Proceedings of the IEEE International Conference on Audio, Speech, and Signal Processing (ICASSP)*, Florence, 2014.
- N. Boulanger-Lewandowski, G. Mysore, M. Hoffman, “Exploiting Long-Term Temporal Dependencies in NMF Using

- Recurrent Neural Networks with Application to Source Separation,” in *Proceedings of the IEEE International Conference on Audio, Speech, and Signal Processing (ICASSP)*, Florence, 2014.
- D. Liang, M. Hoffman, G. Mysore, “A Generative Product-of-Filters Model of Audio,” in *Proceedings of the International Conference on Learning Representations*, Banff, 2014.
- D. Liang, M. Hoffman, D.P.W. Ellis, “Beta Process Sparse Nonnegative Matrix Factorization for Music,” in *Proceedings of the 14th International Conference on Music Information Retrieval (ISMIR)*, Curitiba, 2013.
(Winner, Best Student Paper Award, ISMIR 2009)
- M. Hoffman, C. Wang, J. Paisley, D. Blei, “Stochastic Variational Inference,” in *Journal of Machine Learning Research*, 2013.
- S. Gershman, M. Hoffman, D. Blei, “Nonparametric Variational Inference,” in *Proceedings of the International Conference on Machine Learning*, Edinburgh, 2012.
- D. Mimno, M. Hoffman, D. Blei, “Sparse Stochastic Inference for Latent Dirichlet Allocation,” in *Proceedings of the International Conference on Machine Learning*, Edinburgh, 2012.
- M. Hoffman, “Poisson-Uniform Nonnegative Matrix Factorization,” in *Proceedings of the IEEE Conference on Acoustics, Speech, and Signal Processing*, Kyoto, 2012.
- M. Hoffman, D. Blei, F. Bach, “Online Learning for Latent Dirichlet Allocation,” in *Advances in Neural Information Processing Systems (NIPS)*, Vancouver, 2010. (Spotlight presentation.)
(2nd place Best Student Paper Award, New York Academy of Sciences Machine Learning Symposium 2010)
- M. Hoffman, “Probabilistic Graphical Models for the Analysis and Synthesis of Musical Audio,” Ph.D. Thesis, Princeton University, 2010.
- M. Hoffman, D. Blei, P. Cook, “Bayesian Nonparametric Matrix Factorization for Recorded Music,” in *Proceedings of the International Conference on Machine Learning (ICML)*, Haifa, 2010.
- M. Hoffman, D. Blei, P. Cook, “Easy as CBA: A Simple Probabilistic Model for Tagging Music,” in *Proceedings of the 10th International Conference on Music Information Retrieval (ISMIR)*, Kobe, 2009.
(Winner, Best Student Paper Award, ISMIR 2009)
- M. Hoffman, D. Blei, P. Cook, “Finding Latent Sources in Recorded Music With a Shift-Invariant HDP,” in *Proceedings of the 12th International Conference on Digital Audio Effects*, Como, 2009.
(2nd place Best Student Paper Award, New York Academy of Sciences Machine Learning Symposium 2009)
- M. Hoffman, P. Cook, D. Blei, “Bayesian Spectral Matching: Turning Young MC into MC Hammer via MCMC Sampling,” in *Proceedings of the 2009 International Computer Music Conference*, Montreal, 2009.
- M. Hoffman, D. Blei, P. Cook, “Content-Based Musical Similarity Computation Using the Hierarchical Dirichlet Process,” in *Proceedings of the 9th International Conference on Music Information Retrieval*, Philadelphia, 2008.
- M. Hoffman, P. Cook, “Real-Time Dissonancizers: Two Dissonance-Augmenting Audio Effects,” in *Proceedings of the 11th International Conference on Digital Audio Effects*, Espoo, 2008.
- M. Hoffman, P. Cook, D. Blei, “Data-driven recomposition using the hierarchical Dirichlet process hidden Markov model,” in *Proceedings of the 2008 International Computer Music Conference*, Belfast, 2008.
- M. Hoffman, P. Cook, “The FeatSynth Framework for Feature-Based Synthesis: Design and Applications,” in *Proceedings of the 2007 International Computer Music Conference*, Copenhagen, 2007.
- M. Hoffman, P. Cook, “Real-time Feature-Based Synthesis for Live Musical Performance,” in *Proceedings of the 2007 International Conference on New Interfaces for Musical Expression*, New York, 2007.
- M. Hoffman and P. Cook, “Feature-Based Synthesis: Mapping from Acoustic and Perceptual Features to Synthesis Parameters,” in *Proceedings of the International Computer Music Conference*, New Orleans, 2006.
- Z. Wang, M. Hoffman, P. Cook, K. Li, “VFerret: Content-Based Similarity Tool for Continuous Archived Video,” in *Proceedings of the 3rd ACM Workshop on Continuous Archival and Retrieval of Personal Experiences*, Santa Barbara, 2006.
- M. Hoffman, P. Cook, “Feature-Based Synthesis: A Tool for Evaluating, Designing, and Interacting with Music-IR Systems,” in *Proceedings of the 7th International Conference on Music Information Retrieval*, Victoria, 2006.
- M. Hoffman, P. Cook, D. Vilkomerson, “Staining Doppler Audio,” in *Proceedings of the IEEE International Ultrasonics Symposium*, Vancouver, 2006.
- M. Hoffman, P. Cook, “Feature-Based Synthesis for Sonification and Psychoacoustic Research,” in *Proceedings of the International Conference on Auditory Display*, London, 2006.