

(b) *Prior to coming to UBC*

University, Company or Organization	Rank or Title	Dates
Invrea Ltd.	Director	Jan 2016–present
Alan Turing Institute	Turing Fellow	Feb 2016–Mar 2018
University of Oxford	Associate Professor	Apr 2013–Mar 2018
Kellogg College, University of Oxford	Governing Body Fellow	Apr 2013–Mar 2018
CCLS, Columbia University	Research Scientist	Aug 2012–Dec 2012
Betacular, Ltd.	Founder/Director	Aug 2010–Apr 2018
Columbia University	Assistant Professor	Aug 2009–Aug 2012
Stan James, Ltd.	Consultant	2008–2009
Gatsby Unit, University College London	Postdoctoral Fellow	June 2007–Aug 2009
Interfolio, Inc.	CEO	2002
America Online	Principal Engineer	2001–2002
ToFish!, Inc.	CEO/Founder	1998–2000
Lawrence Berkeley National Laboratory	Research Engineer	1997-1998
Cornell Theory Center	Research Engineer	1996-1997

(c) *At UBC*

Rank or Title	Dates
Associate Professor	April 2018–present

(c) *Date of granting of tenure at U.B.C.:* 16 April 2018

7. LEAVES OF ABSENCE

University, Company or Organization at which Leave was taken	Type of Leave	Dates
---	---------------	-------

8. TEACHING(a) *Areas of special interest and accomplishments*(b) *Courses Taught at UBC*

Session	Course Number	Scheduled Hours	Class Size	Hours Taught			
				Lectures	Tutorials	Labs	Other
Spring 2019	CPSC 340 (** NEW **))	39	173	39	—	—	—
Fall 2018	CPSC 532/539W (** NEW **))	39	21	28	—	—	—

(c) *Graduate Students Supervised*

Student Name	Program Type	Year		Principal Supervisor	CoSupervisor
		Start	Finish		
Wilder Lavington (***) NEW (***)	Ph.D.	18		Wood	Schmidt
Christian Weilbach (***) NEW (***)	Ph.D.	18		Wood	
Will Harvey (***) NEW (***)	Ph.D.	18		Wood	
Boyan Bernov (***) NEW (***)	Ph.D.	18		Wood	
Andreas Munk (***) NEW (***)	Ph.D.	18		Wood	
Vaden Masrani (***) NEW (***)	Ph.D.	17		Wood	
Rob Zinkov	Ph.D. ^o	17		Wood	
Michael Teng	Ph.D. ^o	17		Wood	
Bradley Gram-Hansen	Ph.D. ^o	16		Wood	Teh
Andrew Warrington	Ph.D. ^o	16		Wood	
Max Igl	Ph.D. [†]	16		Wood	Whiteson
Adam Golinski	Ph.D. [†]	16		Wood	Teh
Tuan Anh Le	Ph.D. [†]	15	19	Wood	
Tom Rainforth	Ph.D. [†]	14	17	Wood	Osborne
Brooks Paige	Ph.D. [†]	13	16	Wood	
Saeid Naderiparizi (***) NEW (***)	M.S.	17	19	Wood	
Mario Lezcano Casado	M.S. [†]	16	17	Wood	
David Martinez Rubio	M.S. [†]	16	17	Wood	
Yura Perov	M.S. [†]	14	16	Wood	
Brooks Page	M.S. [*]	11	13	Wood	
David Pfau	M.S. [*]	11	13	Wood	
Nicholas Bartlett	M.S. [*]	11	13	Wood	
Jan Gasthaus	M.S. ⁺	06	07	Wood	
Will Harvey	M.Eng. [†]	16	17	Wood	
Arthur Spencer	M.Eng. [†]	16	17	Wood	
Billy Smith	M.Eng. [†]	15	16	Wood	
Andrew Warrington	M.Eng. [†]	15	16	Wood	
Peter Czaban	M.Eng. [†]	15	16	Wood	
Bo Moon	M.Eng. [†]	15	16	Wood	
Dave Janz	M.Eng. [†]	15	16	Wood	
Tuan Anh Le	M.Eng. [†]	14	15	Wood	
Lawrence Middleton	M.Eng. [†]	13	14	Wood	
Becky Dawes	M.Eng. [†]	13	14	Wood	

⁺ from UCL; [†] from Oxford; ^o from Oxford, UBC VIRS; ^{*} from Columbia

(d) *Postdoctoral Fellows Supervised*

- Adam Scibior, (Ph.D. from Cambridge) February 2019–present. (***) NEW (***)Funded by Startup/DARPA/Inverted AI/Mitacs.
- Gunes Baydin, (Ph.D. from Universitat de Barcelona) September 2016–February 2019. Funded by DARPA and Intel.
- Marcin Szymczak, (Ph.D. from Edinburgh) October 2017–April 2018. Funded by DARPA.
- Tobias Kohn, (Ph.D. from ETH) December 2017–April 2018. Funded by DARPA.
- Jan Willem van de Meent (Ph.D from Leiden) May 2013–July 2016. Funded by DARPA.

- David Tolpin (Ph.D from Ben Gurion) May 2014–August 2015. Funded by DARPA.

- (e) *Continuing Education Activities*
- (f) *Visiting Lecturer indicate university/organization and dates*
- (g) *Other*

9. SCHOLARLY AND PROFESSIONAL ACTIVITIES

- (a) *Areas of special interest and accomplishments*

My primary research areas are probabilistic programming and applied statistical machine learning. My research interests range from the development of new probabilistic models and inference algorithms to real-world applications. My research contributions include probabilistic programming systems, new models and inference algorithms, and novel applications of such models to problems in neuroscience, natural language processing, robotics, and compression.

- (b) *Research or equivalent grants (indicate under COMP whether grants were obtained competitively (C) or non-competitively (NC))*

Granting Agency	Subject	Comp	\$ Per Year	Year	Principal Investigator	Co-Investigators
VPRI CFI	Digital Innovation In Composites Manufacturing (***) ONGOING (***)	C	???	19-24	Poursartip	Wood and many others
MITACS Accelerate / Inverted AI	Probabilistic Programming for Autonomous Driving (***) ONGOING (***)	C	\$400k	19-21	van de Panne	Wood
DARPA	Data Driven Discovery of Models (***) ONGOING (***)	C	\$530k USD	19-23	Wood	
DARPA / CRA	Tractable High-Capacity Probabilistic Models for Learning with Less Labels (***) NEW (***)	C	\$115k	19-21	Wood	Sigal
Compute Canada	General-Purpose Automation of Machine Learning (***) NEW (***)	C	\$480k CAD equiv.	19	Wood	Leyton-Brown
CIFAR	CIFAR AI Chairs Program (***) NEW (***)	C	\$110k USD	18-23	Wood	
NSERC DTA	Advanced Probabilistic Programming (***) NEW (***)	C	\$40k	18-20	Wood	
NSERC Discovery	Advanced Probabilistic Programming (***) NEW (***)	C	\$55k	18-22	Wood	
DARPA	Data Driven Discovery of Models	C	\$450k USD	17-18	Wood	
Intel	Inference Comp. for High Energy Physics	C	\$100k USD	17-19	Wood	
Alan Turing Institute	Probabilistic programming workshop	C	£65k	15	Wood	
Microsoft	Probabilistic programming	NC	£8k	14	Wood	
British Petroleum	Automated pipeline inspection	C	\$100k USD	14-17	Wood	Osborn, Vedaldi
DARPA	Probabilistic programming and advanced machine learning	C	\$300k USD	14-18	Goodman	Wood, Hanrahan
Amazon	Research computing award	C	\$10k USD	14	Wood	
Google	Bayesian nonparametric modeling	C	\$70k USD	14	Wood	
Xerox	Bayesian nonparametric modeling	C	\$90k USD	14	Wood	

(c) *Research or equivalent contracts (indicate under COMP whether grants were obtained competitively (C) or non-competitively (NC))*

(d) *Invited Presentations*

- “Challenges at the confluence of deep learning and probabilistic programming” NIPS Bayesian Deep Learning Workshop, Montreal, QC, 2018¹ (***) **NEW** (***)
- “Inference Compilation or How I Learned to Stop Worrying and Love Deep Networks ” Probabilistic Programming Conference, Boston, MA, 2018 (***) **NEW** (***)
- “Deep Probabilistic Learning and Inference” Element.AI UBC Workshop, Vancouver, BC, 2018
- “Working Towards Distributed Inference Compilation at Scale” LBNL/Intel BDC PCC Workshop, Berkeley, CA, 2018
- “Probabilistic Programming and Inference Compilation, or, How I Learned to Stop Worrying and Love Deep Networks,” CIFAR Learning in Machines and Brains, Long Beach, CA, 2017
- “Deep Probabilistic Programming Inference,” Google Research, Zurich, 2017
- “Probabilistic Programming, Bayesian Nonparametrics, and Inference Compilation” BISP, Milan, 2017
- “Machine Learning and Probabilistic Programming,” British Embassy, Tokyo, Fujitsu, Nagoya Chamber of Commerce, Preferred Networks, Softbank, 2017
- “Revolutionizing Decision Making, Democratizing Data Science, and Automating Machine Learning via Probabilistic Programming,” Loughborough University 2016, NVIDIA, 2017
- “Probabilistic Programming; Ways Forward,” Google, Berkeley, 2015
- “Simulators as priors and neuroscience applications” Janelia Farm Workshop : Big Data Workshop 2015
- “Learning Automata with Infinite State Cardinality,” MIT, 2013
- “Bayesian (Nonparametric) Approaches to Language Modeling,” IBM Watson Research, 2013
- “Bayesian (Nonparametric) Approaches to Language Modeling,” Columbia University, NY, 2012
- “New Bayesian nonparametric tools for statistical machine learning,” University of Illinois at Chicago, City University New York, Oxford, 2012
- “The Infinite Structured Explicit Duration HMM,” ETH, 2012
- “Neuroscience Applications of Dependent Mixtures,” Janelia Farm Workshop : Scaling up EM Connectomics 2012
- “The Infinite Structured Explicit Duration HMM” ISBA, Kyoto, 2012
- “The Sequence Memoizer” Information Theory and Applications, UCSD, 2011
- “Inference in Explicit Duration Hidden Markov Models,” University of Pennsylvania, 2011

¹Average attendance: 1000

- “The Sequence Memoizer,” Columbia University, Brown University, University of Edinburgh, Oxford University, Australia National University 2009; ITA, 2011
- “Nonparametric Bayesian Natural Language Model Domain Adaptation,” Columbia University, Princeton University, University of Utah, 2009
- “Nonparametric Bayesian Natural Language Model Domain Adaptation,” Radboud University, NL and Cambridge University, UK, 2007
- “A Nonparametric Bayesian Alternative to Spike Sorting,” University College London, UK and Radboud University, NL, 2007
- “Gentle Introduction to Infinite Gaussian Mixture Modeling,” Brown University, RI, 2006
- “Bayesian Decoding for Neural Prostheses,” Northwestern University, IL, 2005
- “Variability of Manual Spike Sorting for Multi-Electrode Arrays,” University of Chicago, IL, 2003

(e) *Other Presentations*

(f) *Other*

(g) *Tutorials*

- “Probabilistic Programming,” MLSS Buenos Aires, 2018 (** **NEW** **)
- “Inference Compilation and Universal Probabilistic Programming,” Alan Turing Institute Master Class, Turing Institute, London, 2017
- “Inference Compilation and Universal Probabilistic Programming,” Data on the Brain Video Lecture Series, Berkeley, 2017
- “Inference Compilation and Universal Probabilistic Programming,” Probabilistic Programming Summer School, Portugal, 2017
- “Probabilistic Programming,” DARPA PPAML Summer School, 2016
- “Probabilistic Programming,” Southampton Hackathon, 2016
- “Probabilistic Programming,” NIPS Tutorial², 2015
- “Tutorial on Probabilistic Programming in Machine Learning” Dagstuhl Workshop on “Challenges and Trends in Probabilistic Programming,” 2015
- “Probabilistic Programming,” MLSS Tubingen, 2015
- “Probabilistic Programming,” MLSS Reykjavik, 2014
- “Probabilistic Programming,” Cambridge, 2014
- “Probabilistic Programming,” Imperial, 2014
- “Applied Virtual Reality” SigGraph, Course 14, Los Angeles, CA, 1997

²In the machine learning community being asked to give a tutorial is a major honor. The most significant honor is to be invited to give a NIPS tutorial. The other significant honor is to be asked to teach at the machine learning summer school. There are 6 NIPS tutorials given per year. The audience at mine was over two thousand people. The machine learning summer school (MLSS) series is an exclusive, highly competitive summer school to which very few faculty are invited to lecture (usually approximately 10 per year).

(h) *Conference Participation (Organizer, Keynote Speaker, etc.)*

Keynote Speaker

- “Probabilistic Programming and Inference Compilation or How I Learned to Stop Worrying and Love Deep Networks,” PLDI, Barcelona, 2017
- “Probabilistic Programming,” Inductive Logic Programming, London, 2016
- “Probabilistic Programming,” Artificial General Intelligence, Berlin, 2015
- “Probabilistic Programming; Ways Forward,” DALI, La Palma, 2015

Workshop Organizer

- “NIPS Workshop on Deep Learning for the Physical Sciences,” NIPS, 2017
- “POPL Workshop on Probabilistic Programming Semantics,” POPL, 2016
- “NIPS Workshop on Black Box Learning and Inference,” NIPS, 2015
- “Probabilistic Programming,” Alan Turing Institute, 2015
- “Probabilistic Programming,” DALI, 2015

10. SERVICE TO THE UNIVERSITY(a) *Memberships on committees, including offices held and dates*

- 2018–, CAIDA Steering Committee (***) **NEW** (***)
- 2018–, Huawei SoC-AI Joint Lab Executive Committee (***) **NEW** (***)
- 2018–, Department Faculty Recruiting Committee (***) **NEW** (***)
- 2018–, Department Communications Committee (***) **NEW** (***)
- 2018, Schmidt tenure mini-committee (***) **NEW** (***)

(b) *Other service, including dates*

- Graduate Student Consultative Committee, Oxford Engineering, 2016–2018
- Departmental Video Lecture Capture Coordinator³, 2016–2018
- Kellogg College Finance Committee Fellow⁴ 2014–2018
- Oxford Computer Science Faculty Recruiting Committee, 2016

³This position *introduced* lecture capture in the department of engineering science at Oxford for the first time ever. This involved wiring rooms, selecting equipment, training staff and professors, crafting policy, and conducting beta-tests.

⁴This 6 person-committee reviewed and controlled the finances of the largest graduate college at Oxford on a quarterly basis.

- Oxford Engineering Departmental Foreign Exchange Coordinator⁵, 2013–2015
- Columbia University Statistics Department Computing Committee, 2009
- Gatsby Unit, UCL external talks coordinator, 2008–2009

11. SERVICE TO THE COMMUNITY

(a) *Memberships on scholarly societies, including offices held and dates*

(b) *Memberships on scholarly committees, including offices held and dates*

- NIPS Area Chair⁶, 2011, 2013, 2017, 2019 (***) **NEW** (***)
- ICML Area Chair, 2017, 2019 (***) **NEW** (***)
- IJCAI Senior Program Committee, 2010
- AISTATS Senior Program Committee, 2010, 2013, 2016

(c) *Memberships on other committees, including offices held and dates*

(d) *Editorships (list journal and dates)*

2018– Action Editor, *Journal of Machine Learning Research* (***) **NEW** (***)

(e) *Reviewer (journal, agency, etc. including dates)*

Journal of the Royal Statistical Society, North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Neural Information Processing Systems, Uncertainty in Artificial Intelligence, Artificial Intelligence and Statistics, International Conference on Machine Learning, Journal of Machine Learning Research, Association for the Advancement of Artificial Intelligence, Journal of Neuroscience Methods, IEEE Transactions on Biomedical Engineering, IEEE Transactions on Pattern Analysis and Machine Intelligence, International Joint Conferences on Artificial Intelligence, Journal of Statistics and Computing

(f) *External examiner (indicate university and dates)*

Mehran Kazemi, *Representing and Learning Relations and Properties Under Uncertainty*, PhD. UBC, 2018. (***) **NEW** (***)

Christian Steinrucken, *Lossless Data Compression*, PhD. Cambridge University, 2014.

⁵This insubstantially titled role involved completely rewriting Oxford's *institutional policy* on exchange students and renegotiating existing all exchange program agreements with existing partners, particularly Princeton and National University Singapore, as the existing frameworks were found to fall outside of accreditation guidelines.

⁶Area chair at NIPS and other machine learning conferences is the scientific conference management position one below general chair. Responsibilities include recruiting reviewers and managing the review process for between 20-40 papers.

(g) *Consultant (indicate organization and dates)*

(h) *Other service to the community*

12. AWARDS AND DISTINCTIONS

(a) *Awards for Teaching (indicate name of award, awarding organizations and date)*

(b) *Awards for Scholarship (indicate name of award, awarding organizations and date)*

- AISTATS Best Paper Award, 2009

(c) *Awards for Service (indicate name of award, awarding organizations and date)*

(d) *Other Awards*

- Google faculty research award, 2014
- Xerox faculty research award, 2014
- IMSA Alumni Distinguished Leadership Award, 2011
- National Science Foundation REU Award, Cornell Theory Center, 1994
- Honors College Scholar, University of Illinois at Chicago, 1992

13. OTHER RELEVANT INFORMATION (Maximum One Page)**(a) *Abbreviated Research Statement***

I am a computer scientist and statistician; machine-learning is my research area and artificial-intelligence my inspiration. My contributions span probabilistic programming, inference, unsupervised modeling, density estimation, clustering, Bayesian nonparametrics, reinforcement learning, and related subjects. I publish as a computer scientist; my primary conference communities are NIPS, AISTATS, and ICML. I collaborate widely; my past work includes natural language processing, neuroscience, brain computer interfacing, compression, and medical informatics. Current work touches on programming languages, neural networks, and connectomics. In the more distant past I have worked on graphics, visualization, virtual reality, and vision-based image retrieval.

My group is one of a small handful of world-leaders in the field of probabilistic programming. Probabilistic programming, lying at the intersection of machine learning, statistics, and programming languages, advocates automating inference behind a programming language model specification abstraction layer. It is about designing programming languages and runtimes for the same that “do inference.” I have led the development of several leading probabilistic programming languages including Anglican and probabilistic-C.

Probabilistic programming languages (PPL) are on the cusp of becoming practically useful for expressing and solving a wide-range of model-based statistical reasoning problems. The high-level hypothesis my research tests is that continuing PPL research and development will make it possible for the AI community to rapidly develop key new models for perception, reasoning, and action selection that go beyond what current deep learning systems are thought to be capable of now, focusing in particular on semi- and un-supervised model learning and automatic, efficient probabilistic inference in the same. I draw an analogy between what I aim to achieve with my research and how the development and adoption of programming language tools for automating differentiation arguably has led to the deep learning revolution.

THE UNIVERSITY OF BRITISH COLUMBIA

Publication Record

Date: May 7, 2019

Initials:

SURNAME: Wood

FIRST NAME: Frank

MIDDLE NAME: Donald

1. REFEREED PUBLICATIONS

(a) *Journals*

1. F. Caron, W. Neiswanger, F. Wood, A. Doucet, and M. Davy. Generalized Pólya urn for time-varying Pitman-Yor processes. *JMLR*, 18:1–32, 2017.
2. F. Doshi-Velez, D. Pfau, F. Wood, and N. Roy. Bayesian nonparametric methods for partially-observable reinforcement learning. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 37(2):394–407, 2015.
3. A. Perotte, R. Pivovarov, K. Natarajan, N. Weiskopf, F. Wood, and N. Elhadad. Diagnosis code assignment: models and evaluation metrics. *Journal of the American Medical Informatics Association*, 21(2):231–237, 2014.
4. M. Dewar, C. Wiggins, and F. Wood. Inference in hidden Markov models with explicit state duration distributions. *Signal Processing Letters, IEEE*, 19(4):235–238, 2012.
5. F. Wood, J. Gasthaus, C. Archambeau, L. James, and Y.W. Teh. The sequence memoizer. *Communications of the ACM*, 54(2):91–98, 2011.
6. F. Wood and M. J. Black. A non-parametric Bayesian alternative to spike sorting. *Journal of Neuroscience Methods*, 173:1–12, 2008.
7. D. H. Grollman, O. C. Jenkins, and F. Wood. Discovering natural kinds of robot sensory experiences in unstructured environments. *Journal of Field Robotics*, 23:1077–1089, 2006.
8. F. Wood, M. Fellows, C. Vargas-Irwin, M. J. Black, and J. P. Donoghue. On the variability of manual spike sorting. *IEEE Transactions in Biomedical Engineering*, 51:912–918, 2004.
9. F. Wood, D. Brown, B. Amidon, J. Alferness, B. Joseph, R. E. Gillilan, and C. Faerman. Windowing and telecollaboration for virtual reality with applications to the study of a tropical disease. *IEEE Computer Graphics and Applications*, 16:72–78, 1996.
10. R. E. Gillilan and F. Wood. Visualization, virtual reality, and animation within the data flow model of computing. *Computer Graphics*, 29:55–58, 1995.

(b) *Refereed Conference Proceedings*

1. Yuan Zhou, Bradley J Gram-Hansen, Tobias Kohn, Tom Rainforth, Hongseok Yang, and Frank Wood. Lf-ppl: A low-level first order probabilistic programming language for non-differentiable models. In *Proceedings of the Twentieth International Conference on Artificial Intelligence and Statistics*, 2019. (***) **NEW** (***)

2. Michael Teng and Frank Wood. Bayesian distributed stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 6378–6388, 2018. (***) **NEW** (***)
3. Stefan Webb, Adam Golinski, Rob Zinkov, Siddharth Narayanaswamy, Tom Rainforth, Yee Whye Teh, and Frank Wood. Faithful inversion of generative models for effective amortized inference. In *Advances in Neural Information Processing Systems*, pages 3070–3080, 2018. (***) **NEW** (***)
4. M. Igl, L. Zintgraf, T.A. Le, F. Wood, and S. Whiteson. Deep variational reinforcement learning for POMDPs. In *ICML*, 2018. (***) **NEW** (***)
5. T. Rainforth, A. Kosiorek, T. A. Le, C Maddison, M Igl, F Wood, and Y.W. Teh. Tighter variational bounds are not necessarily better. In *ICML*, 2018. (***) **NEW** (***)
6. T. Rainforth, R. Cornish, H. Yang, and F. Wood. On nesting Monte Carlo estimators. In *ICML*, 2018. (***) **NEW** (***)
7. G. Baydin, R. Cornish, D. Martinez-Rubio, M. Schmidt, and F. Wood. Online learning rate adaptation with hypergradient descent. *ICLR*, 2018. (***) **NEW** (***)
8. T.A. Le, M. Igl, T. Jin, T Rainforth, and F. Wood. Auto-encoding Sequential Monte Carlo. *ICLR*, 2018. (***) **NEW** (***)
9. N. Siddarth, B. Paige, A. Desmaison, J.W. van de Meent, F. Wood, N. Goodman, P. Kohli, and P.H.S Torr. Learning disentangled representations with semi-supervised deep generative models. In *Advances in Neural Information Processing Systems*, 2017.
10. N. Dhir, M. Vakar, M. Wijers, A. Markham, and F. Wood. Interpreting lion behaviour as probabilistic programs. In *Uncertainty in Artificial Intelligence*, 2017.
11. T.A. Le, A.G. Baydin, R. Zinkov, and F. Wood. Using synthetic data to train neural networks is model-based reasoning **oral**. In *30th International Joint Conference on Neural Networks*, 2017.
12. T.A. Le, A.G. Baydin, and F. Wood. Inference Compilation and Universal Probabilistic Programming. In *20th International Conference on Artificial Intelligence and Statistics*, 2017.
13. T. Rainforth, T.A. Le, J.W. van de Meent, and F. Wood. Bayesian optimization for probabilistic programs. In *Advances in Neural Information Processing Systems*, 2016.
14. N. Dhir, Y. Perov, and F. Wood. Nonparametric Bayesian models for unsupervised activity recognition and tracking. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2016)*, 2016.
15. S. Staton, H. Yang, C. Heunen, O. Kammar, and F. Wood. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Thirty-First Annual ACM/IEEE Symposium on Logic In Computer Science*, 2016.
16. Y. Perov and F. Wood. Automatic sampler discovery via probabilistic programming and approximate Bayesian computation. In *Artificial General Intelligence*, pages 262–273, 2016.
17. B. Paige, D. Sejdinovic, and F. Wood. Super-sampling with reservoir. In *Proceedings of the 32nd Annual Conference on Uncertainty in Artificial Intelligence (UAI-2016)*. AUAI Press, 2016.
18. B. Paige and F. Wood. Inference networks for sequential Monte Carlo in graphical models **oral**. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *JMLR*, 2016.

19. T. Rainforth, C.A. Naesseth, F. Lindsten, B. Paige, J.W. van de Meent, A. Doucet, and F. Wood. Interacting particle Markov chain Monte Carlo **oral**. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *JMLR*, 2016.
20. D. Tolpin, J. W. van de Meent, and F. Wood. Probabilistic programming in Anglican - **oral; demo**. In *Machine Learning and Knowledge Discovery in Databases*, pages 308–311. Springer International Publishing, 2015.
21. D. Tolpin, J. W. van de Meent, B. Paige, and F. Wood. Output-sensitive adaptive Metropolis-Hastings for probabilistic programs - **oral**. In *Machine Learning and Knowledge Discovery in Databases*, pages 311–326. Springer International Publishing, 2015.
22. D. Tolpin and F. Wood. Maximum a posteriori estimation by search in probabilistic programs - **oral**. In *Eighth Annual Symposium on Combinatorial Search*, 2015.
23. J. W. van de Meent, H. Yang, V. Mansinghka, and F. Wood. Particle Gibbs with ancestor sampling for probabilistic programs. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pages 986–994, 2015.
24. B. Paige, F. Wood, A. Doucet, and Y.W. Teh. Asynchronous anytime sequential Monte Carlo - **oral**. In *Advances in Neural Information Processing Systems*, pages 3410–3418, 2014.
25. B. Paige and F. Wood. A compilation target for probabilistic programming languages - **oral**. In *JMLR; ICML 2014*, pages 1935–1943, 2014.
26. F. Wood, J. W. van de Meent, and V. Mansinghka. A new approach to probabilistic programming inference - **oral**. In *Artificial Intelligence and Statistics*, pages 1024–1032, 2014.
27. W. Neiswanger, F. Wood, and E. Xing. The dependent Dirichlet process mixture of objects for detection-free tracking and object modeling. In *Artificial Intelligence and Statistics*, pages 660–668, 2014.
28. M. Elsner, S. Goldwater, N. Feldman, and F. Wood. A joint learning model of word segmentation, lexical acquisition, and phonetic variability. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 42–54, 2013.
29. J. W. van de Meent, J. E. Bronson, R. L. Gonzalez Jr., F. Wood, and C. H. Wiggins. Learning biochemical kinetic models from single-molecule data with hierarchically-coupled hidden Markov models. In *International Conference on Machine Learning*, pages 361–369, 2013.
30. C. Smith, F. Wood, and L. Paninski. Low rank continuous-space graphical models - **oral**. In *Artificial Intelligence and Statistics*, volume 152, pages 1064–1072, 2012.
31. A. Perotte, N. Bartlett, N. Elhadad, and F. Wood. Hierarchically supervised latent Dirichlet allocation. In *Advances in Neural Information Processing Systems*, pages 2609–2617, 2011.
32. N. Bartlett and F. Wood. Deplump for streaming data. In *Data Compression Conference*, pages 363–372, 2011.
33. D. Pfau, N. Bartlett, and F. Wood. Probabilistic deterministic infinite automata. In *Advances in Neural Information Processing Systems - spotlight*, pages 1930–1938, 2011.
34. N. Bartlett, D. Pfau, and F. Wood. Forgetting counts : Constant memory inference for a dependent hierarchical Pitman-Yor process - **oral**. In *Proceedings of the 26th International Conference on Machine Learning*, pages 63–70, 2010.

35. J. Gasthaus, F. Wood, and Y.W. Teh. Lossless compression based on the Sequence Memoizer. In *Data Compression Conference*, pages 337–345, 2010.
36. F. Wood, C. Archambeau, J. Gasthaus, L. James, and Y.W. Teh. A stochastic memoizer for sequence data - **oral**. In *Proceedings of the 26th International Conference on Machine Learning*, pages 1129–1136, 2009.
37. F. Wood and Y.W. Teh. A hierarchical nonparametric Bayesian approach to statistical language model domain adaptation - **oral; best paper**. In *Artificial Intelligence and Statistics*, pages 607–614, 2009.
38. J. Gasthaus, F. Wood, D. Görür, and Y.W. Teh. Dependent Dirichlet process spike sorting. In *Advances in Neural Information Processing Systems*, pages 497–504, 2009.
39. P. Berkes, J. W. Pillow, and F. Wood. Characterizing neural dependencies with Poisson copula models. In *Advances in Neural Information Processing Systems*, pages 129 – 136, 2009.
40. F. Wood and T. L. Griffiths. Particle filtering for non-parametric Bayesian matrix factorization. In *Advances in Neural Information Processing Systems*, pages 1513–1520, 2006.
41. F. Wood, S. Goldwater, and M. J. Black. A non-parametric Bayesian approach to spike sorting. In *Proceedings of the 28th IEEE Conference on Engineering in Medicine and Biological Systems*, pages 1165–1169, 2006.
42. F. Wood, T. L. Griffiths, and Z. Ghahramani. A non-parametric Bayesian method for inferring hidden causes. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, pages 536–543, 2006.
43. S. P. Kim, F. Wood, and M. J. Black. Statistical analysis of the non-stationarity of neural population codes. In *The First IEEE / RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics*, pages 259–299, 2006.
44. F. Wood, S. Roth, and M. J. Black. Modeling neural population spiking activity with Gibbs distributions. In *Advances in Neural Information Processing Systems*, pages 1537–1544, 2005.
45. F. Wood, Prabhat, J. P. Donoghue, and M. J. Black. Inferring attentional state and kinematics from motor cortical firing rates. In *Proceedings of the 27th IEEE Conference on Engineering in Medicine and Biological Systems*, pages 149–152, 2005.
46. F. Wood, M. Fellows, J. P. Donoghue, and M. J. Black. Automatic spike sorting for neural decoding - **oral**. In *Proceedings of the 27th IEEE Conference on Engineering in Medicine and Biological Systems*, pages 4126–4129, 2004.

(c) *Other*

2. NON-REFEREED PUBLICATIONS

(a) *Journals*

(b) *Conference and Workshop Proceedings*

1. Mitar Milutinovic, Atılım Güneş Baydin, Robert Zinkov, William Harvey, Dawn Song, Frank Wood, and Wade Shen. End-to-end training of differentiable pipelines across machine learning frameworks. 2017.

2. M. Lezcano-Casado, A.G. Baydin, D. Martinez-Rubio, T.A. Le, F. Wood, G. Heinrich, L. and Louppe, K. Cranmer, K. Ng, W. Bhimji, and Prbhat. Improvements to inference compilation for probabilistic programming in large-scale scientific simulators. In *NeurIPS Workshop on Deep Learning for the Physical Sciences*, 2017.
3. T. Rainforth, T.A. Le, C. Igl, M. amd Maddison, and F. Wood. Tighter variational bounds are not necessarily better. In *NeurIPS Workshop on Bayesian Deep Learning*, 2017.
4. M. Milutinovic, A.G. Baydin, R. Zinkov, W. Harvey, D. Song, F. Wood, and W. Shen. End-to-end training of differentiable pipelines across machine learning frameworks. In *NeurIPS Workshop on Automatic Differentiation*, 2017.
5. R. Cornish, H. Yang, and F. Wood. Towards a testable notion of generalization for generative adversarial networks. In *NeurIPS Workshop on Deep Learning: Bridging Theory and Practice*, 2017.
6. Robert Zinkov Stefan Webb, Adam Golinski and Frank Wood. Principled inference networks in deep generative models. In *NeurIPS Workshop on Bayesian Deep Learning*, 2017.
7. T. Rainforth, Y. Zhou, J.W. van de Meent, H. Yang, and F. Wood. Inference trees: Adaptive inference with exploration. In *NeurIPS Workshop on Bayesian Deep Learning*, 2017.
8. J. W. van de Meent, B. Paige, D. Tolpin, and F. Wood. An interface for black box learning in probabilistic programs. In *POPL Workshop on Probabilistic Programming Semantics*, 2016.
9. S. Staton, H. Yang amd C. Heunen, O. Kammar, and F. Wood. Semantics of higher-order probabilistic programs. In *POPL Workshop on Probabilistic Programming Semantics*, 2016.
10. B. Paige and F. Wood. Inference networks. In *NeurIPS Workshop on Advances in Approximate Bayesian Inference*, 2015.
11. T. Rainforth and F. Wood. Bayesian optimization for probabilistic programs. In *NeurIPS Workshop on Black Box Learning and Inference*, 2015.
12. Y. Perov, T. A. Le, and F. Wood. Data-driven sequential Monte Carlo in probabilistic programming. In *NeurIPS Workshop on Black Box Learning and Inference*, 2015.
13. D. Tolpin, B. Paige, J. W. van de Meent, and F. Wood. Path finding under uncertainty through probabilistic inference. In *In Proceedings of the 5th Workshop on Planning and Learning, ICAPS*, 2015.
14. D. Tolpin, J. W. van de Meent, B. Paige, and F. Wood. Adaptive scheduling in MCMC and probabilistic programming. In *NeurIPS Probabilistic Programming Workshop*, 2014.
15. J. W. van de Meent, H. Yang, V. Mansinghka, and F. Wood. Particle Gibbs with ancestor resampling for probabilistic programs. In *NeurIPS Probabilistic Programming Workshop*, 2014.
16. Y. N. Perov and F. Wood. Learning probabilistic programs. In *NeurIPS Probabilistic Programming Workshop*, 2014.
17. T. B. Paige, X. Zhang, J. Forde, and F. Wood. Perspective inference for eye-to-eye videoconferencing: Empirical evaluation tools and data. In *Proc. 7th Annual Machine Learning Workshop, New York Academy of Sciences*, 2012.
18. F. Wood. Modeling streaming data in the absence of sufficiency. In *NIPS Nonparametric Bayes Workshop*, 2011.

19. F. Wood and Y.W. Teh. A hierarchical, hierarchical Pitman Yor process language model. In *ICML/UAI Nonparametric Bayes Workshop*, 2008.
20. D. H. Grollman, O. C. Jenkins, and F. Wood. Discovering natural kinds of robot sensory experiences in unstructured environments. In *Advances in Neural Information Processing Systems Workshop on Machine Learning Based Robotics in Unstructured Environments*, 2005.

(c) *Other*

1. Michael Teng, Tuan Anh Le, Adam Scibior, and Frank Wood. Imitation learning of factored multi-agent reactive models. *arXiv preprint arXiv:1903.04714*, 2019. (***) **NEW** (***)
2. Atilim Gunes Baydin, Lukas Heinrich, Wahid Bhimji, Bradley Gram-Hansen, Gilles Louppe, Lei Shao, Kyle Cranmer, Frank Wood, et al. Efficient probabilistic inference in the quest for physics beyond the standard model. *arXiv preprint arXiv:1807.07706*, 2018. (***) **NEW** (***)
3. Tom Rainforth, Yuan Zhou, Xiaoyu Lu, Yee Whye Teh, Frank Wood, Hongseok Yang, and Jan-Willem van de Meent. Inference trees: Adaptive inference with exploration. *arXiv preprint arXiv:1806.09550*, 2018. (***) **NEW** (***)
4. Tuan Anh Le, Adam R Kosiorek, N Siddharth, Yee Whye Teh, and Frank Wood. Revisiting reweighted wake-sleep. *arXiv preprint arXiv:1805.10469*, 2018. (***) **NEW** (***)
5. Bradley Gram-Hansen, Yuan Zhou, Tobias Kohn, Hongseok Yang, and Frank Wood. Discontinuous Hamiltonian Monte Carlo for probabilistic programs. *arXiv preprint arXiv:1804.03523*, 2018. (***) **NEW** (***)
6. Michael Teng and Frank Wood. High throughput synchronous distributed stochastic gradient descent. *arXiv preprint arXiv:1803.04209*, 2018. (***) **NEW** (***)
7. Robert Cornish, Hongseok Yang, and Frank Wood. Towards a testable notion of generalization for generative adversarial networks. 2018. (***) **NEW** (***)
8. Mario Lezcano Casado, Atilim Gunes Baydin, David Martínez Rubio, Tuan Anh Le, Frank Wood, Lukas Heinrich, Gilles Louppe, Kyle Cranmer, Karen Ng, Wahid Bhimji, et al. Improvements to inference compilation for probabilistic programming in large-scale scientific simulators. *arXiv preprint arXiv:1712.07901*, 2017.
9. Andrew Warrington and Frank Wood. Updating the vesicle-cnn synapse detector. *arXiv preprint arXiv:1710.11397*, 2017.
10. D. Janz, B. Paige, T. Rainforth, J.W. van de Meent, and F. Wood. Probabilistic structure discovery in time series data. *arXiv preprint arXiv:1611.06863*, 2016.
11. T. Rainforth, R. Cornish, H. Yang, and F. Wood. On the pitfalls of nested Monte Carlo. *arXiv preprint arXiv:1612.00951*, 2016.
12. D. Tolpin, J. W. van de Meent, B. Paige, and F. Wood. Adaptive scheduling in MCMC and probabilistic programming. *arXiv preprint arXiv: 1501.05677*, 2015.
13. T. Rainforth and F. Wood. Canonical correlation forests. *arXiv preprint arXiv:1507.05444*, 2015.
14. Y. Perov and F. Wood. Learning probabilistic programs. *arXiv preprint arXiv:1407.2646*, 2014.

15. Jonathan H Huggins and F. Wood. Infinite structured hidden semi-Markov models. *arXiv preprint arXiv:1407.0044*, 2014.
16. J. W. van de Meent, B. Paige, and F. Wood. Tempering by subsampling. *arXiv preprint arXiv:1401.7145*, 2014.
17. J. Zech and F. Wood. Inferring team strengths using a discrete Markov random field. *arXiv preprint arXiv:1305.1998*, 2013.
18. W. Neiswanger and F. Wood. Unsupervised detection and tracking of arbitrary objects with dependent dirichlet process mixtures. *arXiv preprint arXiv:1210.3288*, 2012.
19. F. Wood, D.H. Grollman, K.A. Heller, O.C. Jenkins, and M.J. Black. Incremental nonparametric Bayesian regression. Technical Report CS-08-07, Brown University, Department of Computer Science, 2008.

3. BOOKS

(a) *Authored*

1. Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, Frank Wood, Introduction to Probabilistic Programming, FTML, NOW Publishers, In Review, 2018. (225 pages)

(b) *Edited*

(c) *Chapters*

1. F. Wood, A. Perotte. Mixed Membership Classification for Documents with Hierarchically Structured Labels. in *Handbook on Mixed Membership Models*. Ed. Airoldi EM, Blei D, Erosheva EA, and Fienberg E. Chapman and Hall/CRC 2014.

4. PATENTS

1. F. Wood, H. Yang, M. Wu, Y. Perov, *Computing Engine, Software, System and Method* (application), 2016
2. F. Wood, B. Paige, Y. Perov, J.W. van de Meent, V. Mansinghka, *Methods and devices for executing program code of a probabilistic programming language* (application), 2015
3. M. J. Black, W. Wu, and F. Wood, 11,086,956, *Method and system for automatic decoding of motor cortical activity*, 2005
4. G. Pass and F. Wood, 6,671,402, *Representing an image with weighted joint histogram*, 2003
5. G. Pass and F. Wood, 6,522,782, *Image and text searching techniques*, 2003
6. G. Pass and F. Wood, 6,556,710, *Image searching techniques*, 2003
7. G. Pass and F. Wood, 6,622,780, *Indexing of images and/or text*, 2003
8. G. Pass and F. Wood, 6,522,779, *Representing an image with a posterized joint histogram*, 2003

5. SPECIAL COPYRIGHTS**6. ARTISTIC WORKS, PERFORMANCES, DESIGNS****(a) Software**

- **Anglican**⁷ is an open source, compiled probabilistic programming language integrated with Clojure, a general purpose functional programming language that just-in-time compiles to the Java Virtual Machine (JVM). It is one of the most widely downloaded higher-order probabilistic programming languages in the world. It has been used for instruction at universities in Canada, Korea, and the UK. It also forms the software foundation for several emerging AI startups. I wrote the first version of this language and have since then been responsible for resourcing and guiding its continued development.
- **PyProb**⁸ (***) is a PyTorch-based library for probabilistic programming and inference compilation. The main focus of this library is on coupling existing simulation codebases with probabilistic inference with minimal intervention.

7. OTHER WORKS**8. WORK SUBMITTED (including publisher and date of submission)****9. WORK IN PROGRESS (including degree of completion)**

⁷<https://anglican.ml>

⁸<https://github.com/probprog/pyprob>