

**THE UNIVERSITY OF BRITISH COLUMBIA**

*Curriculum Vitae for Faculty Members*

**Date:** October 15, 2022

Initials:

1. **SURNAME:** Wood

**FIRST NAME:** Frank

**MIDDLE NAME:** Donald

2. **DEPARTMENT/SCHOOL:** Computer Science

3. **FACULTY:** Science

4. **PRESENT RANK:** Associate Professor

**SINCE:** 16 April 2018

5. **POST-SECONDARY EDUCATION**

(a) *Degree*

University or Institution	Degree	Subject Area	Dates
Cornell University	B.S.	Comp. Sci.	1996
Brown University	M.S.	Comp. Sci.	2004
Brown University	Ph.D.	Comp. Sci.	2007

(b) *Special Professional Qualifications*

6. **EMPLOYMENT RECORD**

(a) *While at UBC*

University, Company or Organization	Rank or Title	Dates
DeRisk Ltd.	Director	Apr 2022–present
University of British Columbia	Special Advisor, Res. and Ind. Partner.	Mar 2021–present
MILA	Canada CIFAR AI Chair	May 2019–present
Inverted AI Ltd.	Director	Nov 2018–present

(b) *Prior to coming to UBC*

University, Company or Organization	Rank or Title	Dates
Invrea Ltd.	Director	Jan 2016–Jan 2021
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
Fall 2022	CPSC 532	39	4	28	—	—	—
Fall 2021	CPSC 532	39	18	28	—	—	—
Spring 2021	CPSC 532	39	21	28	—	—	—
Winter 2020	CPSC 340	39	144	39	—	—	—
Winter 2020	CPSC 340	39	137	39	—	—	—
Spring 2020	CPSC 340	39	165	39	—	—	—
Spring 2019	CPSC 340	39	173	39	—	—	—
Fall 2018	CPSC 532/539W	39	21	28	—	—	—

(c) *Graduate Students Supervised*

Student Name	Program Type	Year		Principal Supervisor	CoSupervisor
		Start	Finish		
Dylan Green	M.S.	22		Wood	Schmidt
Xiaoxuan Liang	Ph.D.	22		Wood	
Matthew Niedoba	Ph.D.-Track MSc Student	22		Wood	
Larry Liu	Ph.D.	21		Wood	
Justice Sefas	Ph.D.-Track MSc Student	21		Wood	
Vasileios Lioutas	Ph.D.	20		Wood	
Jason Yoo	Ph.D.	20		Wood	
Ryan Fayyazi	Ph.D.-Track MSc Student	20		Wood	
Saeid Naderiparizi	Ph.D.	20		Wood	
Peyman Bateni	M.S.	19	21	Wood	
Wilder Lavington	Ph.D.	18		Wood	Schmidt
Christian Weilbach	Ph.D.	18		Wood	
Will Harvey	Ph.D.	18		Wood	
Andreas Munk	Ph.D.	18		Wood	
Vaden Masrani	Ph.D.	17		Wood	
Michael Teng	Ph.D. <sup>o</sup>	17	21	Wood	
Andrew Warrington	Ph.D. <sup>o</sup>	16	21	Wood	
Rob Zinkov	Ph.D. <sup>o</sup>	17		Wood	
Bradley Gram-Hansen	Ph.D. <sup>o</sup>	16	19	Wood	Teh
Max Igl	Ph.D. <sup>†</sup>	16	19	Wood	Whiteson
Adam Golinski	Ph.D. <sup>†</sup>	16	19	Wood	Teh
Tuan Anh Le	Ph.D. <sup>†</sup>	15	19	Wood	
Tom Rainforth	Ph.D. <sup>†</sup>	14	17	Wood	Osborne
Brooks Paige	Ph.D. <sup>†</sup>	13	16	Wood	
Saeid Naderiparizi	M.S.	17	19	Wood	
Mario Lezcano Casado	M.S. <sup>†</sup>	16	17	Wood	
David Martinez Rubio	M.S. <sup>†</sup>	16	17	Wood	
Yura Perov	M.S. <sup>†</sup>	14	16	Wood	
Brooks Page	M.S. <sup>*</sup>	11	13	Wood	
David Pfau	M.S. <sup>*</sup>	11	13	Wood	
Nicholas Bartlett	M.S. <sup>*</sup>	11	13	Wood	
Jan Gasthaus	M.S. <sup>+</sup>	06	07	Wood	
Will Harvey	M.Eng. <sup>†</sup>	16	17	Wood	
Arthur Spencer	M.Eng. <sup>†</sup>	16	17	Wood	
Billy Smith	M.Eng. <sup>†</sup>	15	16	Wood	
Andrew Warrington	M.Eng. <sup>†</sup>	15	16	Wood	
Peter Czaban	M.Eng. <sup>†</sup>	15	16	Wood	
Bo Moon	M.Eng. <sup>†</sup>	15	16	Wood	
Dave Janz	M.Eng. <sup>†</sup>	15	16	Wood	
Tuan Anh Le	M.Eng. <sup>†</sup>	14	15	Wood	
Lawrence Middleton	M.Eng. <sup>†</sup>	13	14	Wood	
Becky Dawes	M.Eng. <sup>†</sup>	13	14	Wood	

<sup>+</sup> from UCL; <sup>†</sup> from Oxford; <sup>o</sup> from Oxford, *UBC VIRS*; <sup>\*</sup>from Columbia

(d) *Postdoctoral Fellows and Research Associates Supervised*

- Alexander Mead<sup>°</sup>, (Ph.D. from Univ. of Edinburgh) May 2022–present. Funded by LNBL.
- Berend Zwartsenberg<sup>°</sup>, (Ph.D. from UBC) January 2021–January 2022. Funded by DSI-CRN.
- Adam Scibior<sup>°</sup>, (Ph.D. from Cambridge) February 2019–January 2022. Funded by Startup/DARPA/Inverted AI/Mitacs.
- Gunes Baydin<sup>†</sup>, (Ph.D. from Universitat de Barcelona) September 2016–February 2019. Funded by DARPA and Intel.
- Marcin Szymczak<sup>†</sup>, (Ph.D. from Edinburgh) October 2017–April 2018. Funded by DARPA.
- Tobias Kohn<sup>†</sup>, (Ph.D. from ETH) December 2017–April 2018. Funded by DARPA.
- Jan Willem van de Meent<sup>†</sup>, (Ph.D from Leiden) May 2013–July 2016. Funded by DARPA.
- David Tolpin<sup>†</sup>, (Ph.D from Ben Gurion) May 2014–August 2015. Funded by DARPA.

<sup>†</sup> from Oxford; <sup>°</sup> from UBC;

(e) *Continuing Education Activities*

(f) *Visiting Lecturer indicate university/organization and dates*

(g) *Other*

- Philippe Solodov, Undergraduate Research Assistant (under Directed Studies CPSC 448), May 2019 - May 2021.
- Elizabeth Hnatiuk, Undergraduate Research Assistant (under Directed Studies COGS 402), Sep - Dec 2020.
- Olga Solodova, Undergraduate Research Assistant (under Directed Studies COGS 402 and CPSC 448B), Sep 2019 - Aug 2020.
- Jason Yoo, Undergraduate Research Assistant (under Directed Studies COGS 402 and CPSC 448B), Sep 2019 - Aug 2020.
- Youssef Farag, Undergraduate research assistant (under Directed Studies COGS 402 and CPSC 448B ), Sep 2019 - Aug 2020.
- Ryan Fayyazi, Undergraduate Research Assistant (CPSC 449 Honours Theses), Sep 2019 - Aug 2020.
- Onur Tuna, Undergraduate Research Assistant (CPSC 449 Honours Theses), Sep 2019 - Aug 2020.
- Dylan Yung, Undergraduate Research Volunteer, Sep 2019 - Aug 2020.
- Gwen Li, Undergraduate Research Volunteer, Sep 2019 - Mar 2020.
- Ren Wang, Undergraduate Research Assistant, May 2019 - Feb 2020.
- Neil Dhir, Undergraduate summer intern, summer 2019.
- Ryan Fayyazi, Undergraduate summer intern, summer 2019.
- Onur Tuna, Undergraduate summer intern, summer 2019.

- Olga Solodova, Undergraduate summer intern, summer 2019.
- Jason Yoo, Undergraduate summer intern, summer 2019.
- Youssef Farag, Undergraduate summer intern, summer 2019.
- Ray Ding, Undergraduate summer intern, summer 2019.
- Zikun Chen, Undergraduate summer intern, summer 2019.
- Alexander Bergholm, Undergraduate summer intern, summer 2019.
- Elizabeth Hnatiuk, Undergraduate summer intern, summer 2019.

## 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
Compute Canada	Deep Generative Behaviour Modelling	C	\$106k equiv.	22	Wood	Van de Panne
Compute Canada	Learning Expensive-to-Evaluate Policies	C	\$355k equiv.	22	Leyton-Brown	Wood
Department of Energy/LBNL	Surrogating High Dimensional Probability Distributions with Deep Learning for Scientific Inference and Data Analysis	C	\$88k USD	21-24	Wood	Seljak
NSERC Research Tools and Instruments	UBC ML Computational Cluster	C	\$150k	21-22	Wood	Sigal, Schmidt, Leyton-Brown, Campbell
Amazon Research Awards	Automated Machine Learning for Tabular Datasets using Hyperband Embedded Reinforcement Learning	C	\$57k USD + \$100k USD compute	21-22	Leyton-Brown	Wood
Compute Canada	General-Purpose Automation of Machine Learning	C	\$410k equiv.	21	Leyton-Brown	Wood
DND IDEAS	An AI-based Tool for the Analysis of Public Health Policy Options in Response to COVID-19 and other Infectious Diseases	C	\$199k	20-21	Wood	
NSERC Alliance	Data Science & Composite Materials Manufacturing	C	\$300k	20-22	Ng	Wood and many others
CIFAR	AI and COVID-19 Catalyst Grants	C	\$13k	20	Wood	Bloem-Reddy and others
Faculty of Science	Support for Teams to Advance Interdisciplinary Research (STAIR) Grant	C	\$22.5k	20-21	Bloem-Reddy	Wood
Compute Canada	General-Purpose Automation of Machine Learning	C	\$480k equiv.	20	Leyton-Brown	Wood
VPRI CFI	Digital Innovation In Composites Manufacturing	C	\$300k	19-24	Poursartip	Wood and many others
MITACS Accelerate / Inverted AI	Probabilistic Programming for Autonomous Driving	C	\$240k	19-21	Wood	

Granting Agency	Subject	Comp	\$ Per Year	Year	Principal Investigator	Co-Investigators
DARPA	Data Driven Discovery of Models	C	\$700k USD	19-21	Wood	Leyton-Brown, Ligett
DARPA / CRA	Tractable High-Capacity Probabilistic Models for Learning with Less Labels	C	\$140k USD	19-21	Wood	Sigal
Compute Canada	General-Purpose Automation of Machine Learning	C	\$480k equiv.	19	Leyton-Brown	Wood
CIFAR	CIFAR AI Chairs Program	C	\$110k	18-23	Wood	
NSERC DTA	Advanced Probabilistic Programming	C	\$40k	18-20	Wood	
NSERC Discovery	Advanced Probabilistic Programming	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*

- “What is Deep Probabilistic Programming?” Huawei-UBC Workshop, virtual talk, 2021
- “Uncertainty in Artificial Intelligence: Techniques from the Intersection of Deep Learning and Probabilistic Programming” Department of Statistics’ Student Seminar (virtual), University of Michigan 2020
- “Etalumis: Bringing Probabilistic Programming to Scientific Simulators at Scale” Intel HPC



Developer Conference, Denver CO, 2019

- “Challenges at the confluence of deep learning and probabilistic programming” NeurIPS Bayesian Deep Learning Workshop, Montreal, QC, 2018<sup>1</sup>
- “Inference Compilation or How I Learned to Stop Worrying and Love Deep Networks ” Probabilistic Programming Conference, Boston, MA, 2018
- “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
- “The Sequence Memoizer,” Columbia University, Brown University, University of Edinburgh, Oxford University, Australia National University 2009; ITA, 2011

---

<sup>1</sup>Average attendance: 1000

- “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
- “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,” NeurIPS Tutorial<sup>2</sup>, 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

---

<sup>2</sup>In 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 NeurIPS tutorial. The other significant honor is to be asked to teach at the machine learning summer school. There are 6 NeurIPS 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.)*

## Panel Discussion

- “Deploying AI in BC: Challenges and Opportunities”, with Kal Ruberg, Greg Mori, Matt Taddy, Evgueni Loukipoudis, Aline Talhouk, Douglas Kingsford at Emerging Technologies: BC’s AI Showcase, virtual CAIDA conference, Dec 2020
- “How can you use AI to contribute to growing our sustainable and diverse community?”, with Cindy Gordon, Tony Khoo, Amit Varma, Zoe Cayetano and Linus Sebastian at 2020 AI4Youth Canada 3rd Annual National Conference, Aug 2020
- “The future of AI”, with Gary Marcus, Chen Greif, Kevin Leyton-Brown at UBC Robson Square, Feb 2020

## Keynote Speaker

- “Deep Probabilistic Programming: Case Studies and Industrial Opportunities” Emerging Technologies: BC’s AI Showcase, virtual CAIDA conference, 2020
- “Beyond Deep Learning” CrossOver AI, Vancouver, BC, 2019
- “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

- “NeurIPS Workshop on Deep Learning for the Physical Sciences,” NeurIPS, 2017
- “POPL Workshop on Probabilistic Programming Semantics,” POPL, 2016
- “NeurIPS Workshop on Black Box Learning and Inference,” NeurIPS, 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*

- 2021–, Department Research and Industrial Partnerships Committee
- 2018–, CAIDA Steering Committee
- 2018–, Huawei SoC-AI Joint Lab Executive Committee
- 2018–, Department Faculty Recruiting Committee

- 2020–2021, CRC Tier 2 in Health-Related Areas Recruiting Committee
- 2018–2020, Department Communications Committee
- 2018, Schmidt tenure mini-committee

(b) *Other service, including dates*

- Graduate Student Consultative Committee, Oxford Engineering, 2016–2018
- Departmental Video Lecture Capture Coordinator<sup>3</sup>, 2016–2018
- Kellogg College Finance Committee Fellow<sup>4</sup> 2014–2018
- Oxford Computer Science Faculty Recruiting Committee, 2016
- Oxford Engineering Departmental Foreign Exchange Coordinator<sup>5</sup>, 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*

- Canadian Artificial Intelligence Association (CAIAC) Best Thesis Award Committee, 2021
- AAAI Senior Area Chair, 2020
- NeurIPS Area Chair<sup>6</sup>, 2011, 2013, 2017, 2019
- ICML Area Chair, 2017, 2019
- 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*

---

<sup>3</sup>This 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.

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

<sup>5</sup>This 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.

<sup>6</sup>Area chair at NeurIPS 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.

(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)*

- Ayub Ahmed Gubran, *Models and Techniques for Designing Mobile System-on-Chip Devices*, PhD. UBC, 2020.
- Mehran Kazemi, *Representing and Learning Relations and Properties Under Uncertainty*, PhD. UBC, 2018.
- Christian Steinrucken, *Lossless Data Compression*, PhD. Cambridge University, 2014.

(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)*

- Best Paper Award CVPR Autonomous Driving Workshop “Imagining The Road Ahead: Multi-Agent Trajectory Prediction via Differentiable Simulation,” 2021
- Best Paper Finalist at Supercomputing for “Etalumis: Bringing Probabilistic Programming to Scientific Simulators,” 2019

- ICML Best Paper Honourable Mention for “Amortized Monte Carlo Integration,” 2019
- 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 NeurIPS, 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.

**(b) *Publication Culture***

Computer science as a field is largely, culturally, a conference publication community; machine learning (ML) and artificial intelligence (AI) particularly so. The main conferences in ML/AI are, arguably, NeurIPS, ICML, UAI, AISTATS, AAAI, and ICLR. Acceptance rates at these conferences hover around 20% and submissions are subject to rigorous peer review. Depending on year and conference, usually only an extremely small percentage of papers are given oral presentation slots.

**THE UNIVERSITY OF BRITISH COLUMBIA**  
*Publication Record*

**Date:** October 15, 2022

**Initials:**

**SURNAME:** Wood

**FIRST NAME:** Frank

**MIDDLE NAME:** Donald

**1. REFEREED PUBLICATIONS**

(a) *Journals*

1. Vasileios Lioutas, Adam Ścibior, and Frank Wood. TITRATED: Learned human driving behavior without infractions via amortized inference. *Transactions in Machine Learning Research (TMLR)*, 2022
2. Frank Wood, Andrew Warrington, Saeid Naderiparizi, Christian Weilbach, Vaden Masrani, William Harvey, Adam Ścibior, Boyan Beronov, John Grefenstette, Duncan Campbell, and S. Ali Nasserri. Planning as inference in epidemiological dynamics models. *Frontiers in Artificial Intelligence*, 4, 2022
3. Tom Rainforth, Adam Golinski, Frank Wood, and Sheheryar Zaidi. Target-aware Bayesian inference: How to beat optimal conventional estimators. *Journal of Machine Learning Research*, 21(88):1–54, 2020
4. 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
5. 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
6. 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
7. 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
8. F. Wood, J. Gasthaus, C. Archambeau, L. James, and Y.W. Teh. The sequence memoizer. *Communications of the ACM*, 54(2):91–98, 2011
9. F. Wood and M. J. Black. A non-parametric Bayesian alternative to spike sorting. *Journal of Neuroscience Methods*, 173:1–12, 2008
10. 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
11. 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
12. 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



13. 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. William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weilbach, and Frank Wood. Flexible diffusion modeling of long videos. In *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)*, 2022
2. Jason Yoo and Frank Wood. BayesPCN: A continually learnable predictive coding associative memory. In *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)*, 2022
3. Yunpeng Liu, Jonathan Wilder Lavington, Adam Ścibior, and Frank Wood. Vehicle type specific waypoint generation **oral**. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022
4. Andreas Munk, Berend Zwartsenberg, Adam Ścibior, Atilim Gunes Baydin, Andrew Lawrence Stewart, Goran Fernlund, Anoush Poursartip, and Frank Wood. Probabilistic surrogate networks for simulators with unbounded randomness. In *The 38th Conference on Uncertainty in Artificial Intelligence*, 2022
5. Saeid Naderiparizi, Adam Ścibior, Andreas Munk, Mehrdad Ghadiri, Atilim Gunes Baydin, Bradley J Gram-Hansen, Christian A Schroeder De Witt, Robert Zinkov, Philip Torr, Tom Rainforth, et al. Amortized rejection sampling in universal probabilistic programming - **oral**. In *International Conference on Artificial Intelligence and Statistics*, pages 8392–8412. PMLR, 2022
6. Peyman Bateni, Jarred Barber, Jan-Willem van de Meent, and Frank Wood. Enhancing few-shot image classification with unlabelled examples. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2796–2805, January 2022
7. Vaden Masrani, Rob Brekelmans, Thang Bui, Frank Nielsen, Aram Galstyan, Greg Ver Steeg, and Frank Wood. q-paths: Generalizing the geometric annealing path using power means. In Cassio de Campos and Marloes H. Maathuis, editors, *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Proceedings of Machine Learning Research*, pages 1938–1947. PMLR, 27–30 Jul 2021
8. William Harvey, Saeid Naderiparizi, and Frank Wood. Conditional image generation by conditioning variational auto-encoders. In *International Conference on Learning Representations*, 2022
9. Adam Ścibior, Vaden Masrani, and Frank Wood. Differentiable particle filtering without modifying the forward pass. In *International Conference on Probabilistic Programming (PROBPROG)*, 2021
10. Adam Ścibior, Vasileios Lioutas, Daniele Reda, Peyman Bateni, and Frank Wood. Imagining The Road Ahead: Multi-agent trajectory prediction via differentiable simulation. In *IEEE Intelligent Transportation Systems Conference (ITSC)*, 2021
11. Andreas Munk, William Harvey, and Frank Wood. Assisting the adversary to improve GAN training. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2021. doi: 10.1109/IJCNN52387.2021.9533449
12. Vaden Masrani, Rob Brekelmans, Thang Bui, Frank Nielsen, Aram Galstyan, Greg Ver Steeg, and Frank Wood. q-Paths: Generalizing the geometric annealing path using power means. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2021

13. Boyan Beronov, Christian Weilbach, Frank Wood, and Trevor Campbell. Sequential core-set Monte Carlo. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2021
14. Andreas Munk, William Harvey, and Frank Wood. Assisting the adversary to improve GAN training. In *The International Joint Conference on Neural Networks (IJCNN)*, 2021
15. Andrew Warrington, Jonathan W Lavington, Adam Ścibior, Mark Schmidt, and Frank Wood. Robust asymmetric learning in POMDPs **oral**. In *International Conference on Machine Learning (ICML)*, pages 11013–11023. PMLR, 2021
16. Vu Nguyen, Vaden Masrani, Rob Brekelmans, Michael Osborne, and Frank Wood. Gaussian process bandit optimization of the thermodynamic variational objective. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020
17. Tuan Anh Le, Adam R. Kosiorek, N. Siddharth, Yee Whye Teh, and Frank Wood. Revisiting reweighted wake-sleep for models with stochastic control flow. In Ryan P. Adams and Vibhav Gogate, editors, *Uncertainty in Artificial Intelligence (UAI)*, volume 115 of *Proceedings of Machine Learning Research*, pages 1039–1049, Tel Aviv, Israel, 22–25 Jul 2020. PMLR
18. Michael Teng, Tuan Anh Le, Adam Ścibior, and Frank Wood. Semi-supervised sequential generative models. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2020
19. Christian Weilbach, Boyan Beronov, Frank Wood, and William Harvey. Structured conditional continuous normalizing flows for efficient amortized inference in graphical models. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 4441–4451, 2020
20. Saeid Naderiparizi, Adam Ścibior, Andreas Munk, Mehrdad Ghadiri, Atılım Güneş Baydin, Bradley Gram-Hansen, Christian Schroeder de Witt, Robert Zinkov, Philip HS Torr, Tom Rainforth, et al. Amortized rejection sampling in universal probabilistic programming. In *International Conference on Probabilistic Programming (PROBPROG)*, 2020
21. Rob Brekelmans, Vaden Masrani, Frank Wood, Greg Ver Steeg, and Aram Galstyan. All in the exponential family: Bregman duality in thermodynamic variational inference. In *Thirty-seventh International Conference on Machine Learning (ICML)*, July 2020
22. A Warrington, S Naderiparizi, and F Wood. Coping with simulators that don’t always return. In *The 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020
23. W Harvey, A Munk, AG Baydin, A Bergholm, and F Wood. Attention for inference compilation. In *The second International Conference on Probabilistic Programming (PROBPROG)*, 2020
24. Andreas Munk, Adam Ścibior, AG Baydin, A Stewart, A Fernlund, A Poursartip, and Frank Wood. Deep probabilistic surrogate networks for universal simulator approximation. In *The second International Conference on Probabilistic Programming (PROBPROG)*, 2020
25. Christian Weilbach, Boyan Beronov, Frank Wood, and William Harvey. Structured conditional continuous normalizing flows for efficient amortized inference in graphical models. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 4441–4451, 2020
26. Vaden Masrani, Tuan Anh Le, and Frank Wood. The thermodynamic variational objective. In *Thirty-third Conference on Neural Information Processing Systems (NeurIPS)*, 2019

27. Atılım Güneş Baydin, Lei Shao, Wahid Bhimji, Lukas Heinrich, Lawrence Meadows, Jialin Liu, Andreas Munk, Saeid Naderiparizi, Bradley Gram-Hansen, Gilles Louppe, et al. Etalumis: Bringing probabilistic programming to scientific simulators at scale. In *the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '19)*
28. Adam Goliński, Frank Wood, and Tom Rainforth. Amortized monte carlo integration. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019
29. Tuan Anh Le, Adam R Kosiorek, N Siddharth, Yee Whye Teh, and Frank Wood. Revisiting reweighted wake-sleep for models with stochastic control flow. In *Proceedings of the conference on Uncertainty in Artificial Intelligence (UAI)*, 2019
30. 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
31. Michael Teng and Frank Wood. Bayesian distributed stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 6378–6388, 2018
32. 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
33. M. Igl, L. Zintgraf, T.A. Le, F. Wood, and S. Whiteson. Deep variational reinforcement learning for POMDPs. In *ICML*, 2018
34. 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
35. T. Rainforth, R. Cornish, H. Yang, and F. Wood. On nesting Monte Carlo estimators. In *ICML*, 2018
36. G. Baydin, R. Cornish, D. Martinez-Rubio, M. Schmidt, and F. Wood. Online learning rate adaptation with hypergradient descent. *ICLR*, 2018
37. T.A. Le, M. Igl, T. Jin, T Rainforth, and F. Wood. Auto-encoding Sequential Monte Carlo. *ICLR*, 2018
38. N. Siddharth, 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
39. N. Dhir, M. Vakar, M. Wijers, A. Markham, and F. Wood. Interpreting lion behaviour as probabilistic programs. In *Uncertainty in Artificial Intelligence*, 2017
40. 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
41. 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
42. 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
43. 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

44. 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
45. Y. Perov and F. Wood. Automatic sampler discovery via probabilistic programming and approximate Bayesian computation. In *Artificial General Intelligence*, pages 262–273, 2016
46. 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
47. 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
48. 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
49. 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
50. 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
51. D. Tolpin and F. Wood. Maximum a posteriori estimation by search in probabilistic programs - **oral**. In *Eighth Annual Symposium on Combinatorial Search*, 2015
52. 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
53. 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
54. N. Dhir and F. Wood. Improved activity recognition via kalman smoothing and multiclass linear discriminant analysis. In *Proceedings of the 36th IEEE Conference on Engineering in Medicine and Biological Systems*, 2014
55. B. Paige and F. Wood. A compilation target for probabilistic programming languages - **oral**. In *JMLR; ICML 2014*, pages 1935–1943, 2014
56. 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
57. 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
58. 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

59. 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
60. 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
61. 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
62. N. Bartlett and F. Wood. Deplump for streaming data. In *Data Compression Conference*, pages 363–372, 2011
63. D. Pfau, N. Bartlett, and F. Wood. Probabilistic deterministic infinite automata. In *Advances in Neural Information Processing Systems - spotlight*, pages 1930–1938, 2011
64. 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
65. J. Gasthaus, F. Wood, and Y.W. Teh. Lossless compression based on the Sequence Memoizer. In *Data Compression Conference*, pages 337–345, 2010
66. 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
67. 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
68. 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
69. 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
70. 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
71. 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
72. 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
73. 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 Biomechanics*, pages 259–299, 2006
74. 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

75. 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
76. 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) *Refereed Workshop Proceedings*

1. Adam Ścibior, Vasileios Lioutas, Daniele Reda, Peyman Bateni, and Frank Wood. Imagining The Road Ahead: Multi-agent trajectory prediction via differentiable simulation - **oral; best paper**. In *CVPR Workshop Workshop on Autonomous Driving: Perception, Prediction and Planning*, 2021
2. Jonathan Wilder Lavington, Michael Teng, Mark Schmidt, and Frank Wood. A closer look at gradient estimators with reinforcement learning as inference. In *Deep RL Workshop NeurIPS 2021*, 2021
3. William Harvey, Michael Teng, and Frank Wood. Near-optimal glimpse sequences for improved hard attention neural network training. In *NeurIPS Workshop on Bayesian Deep Learning*, 2019
4. Andrew Warrington, Arthur Spencer, and Frank Wood. The virtual patch clamp: Imputing *C. elegans* membrane potentials from calcium imaging. In *NeurIPS 2019 Workshop Neuro AI homepage*
5. A Warrington, S Naderiparizi, and F Wood. Coping with simulators that don't always return. In *2nd Symposium on Advances in Approximate Bayesian Inference (AABI)*, 2019
6. B Gram-Hansen, C Schroeder de Witt, R Zinkov, S Naderiparizi, A Ścibior, A Munk, F Wood, M Ghadiri, P Torr, Y Whye Teh, A Gunes Baydin, and T Rainforth. Efficient Bayesian inference for nested simulators. In *2nd Symposium on Advances in Approximate Bayesian Inference (AABI)*, 2019
7. 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
8. 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
9. 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
10. 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
11. 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
12. Robert Zinkov Stefan Webb, Adam Golinski and Frank Wood. Principled inference networks in deep generative models. In *NeurIPS Workshop on Bayesian Deep Learning*, 2017
13. 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

14. 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
15. S. Staton, H. Yang and C. Heunen, O. Kammar, and F. Wood. Semantics of higher-order probabilistic programs. In *POPL Workshop on Probabilistic Programming Semantics*, 2016
16. B. Paige and F. Wood. Inference networks. In *NeurIPS Workshop on Advances in Approximate Bayesian Inference*, 2015
17. T. Rainforth and F. Wood. Bayesian optimization for probabilistic programs. In *NeurIPS Workshop on Black Box Learning and Inference*, 2015
18. 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
19. 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
20. 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
21. 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
22. Y. N. Perov and F. Wood. Learning probabilistic programs. In *NeurIPS Probabilistic Programming Workshop*, 2014
23. 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
24. F. Wood. Modeling streaming data in the absence of sufficiency. In *NIPS Nonparametric Bayes Workshop*, 2011
25. F. Wood and Y.W. Teh. A hierarchical, hierarchical Pitman Yor process language model. In *ICML/UAI Nonparametric Bayes Workshop*, 2008
26. 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

(d) *Other*

## 2. NON-REFEREED PUBLICATIONS

(a) *Journals*

(b) *Conference and Workshop Proceedings*

(c) *Other*

1. Berend Zwartsenberg, Adam Ścibior, Matthew Niedoba, Vasileios Lioutas, Yunpeng Liu, Justice Sefas, Setareh Dabiri, Jonathan Wilder Lavington, Trevor Campbell, and Frank Wood. Conditional permutation invariant flows. *arXiv preprint arXiv:2206.09021*, 2022
2. Vasileios Lioutas, Jonathan Wilder Lavington, Justice Sefas, Matthew Niedoba, Yunpeng Liu, Berend Zwartsenberg, Setareh Dabiri, Frank Wood, and Adam Ścibior. Critic sequential monte carlo. 2022
3. William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weilbach, and Frank Wood. Flexible diffusion modeling of long videos. In *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)*, 2022
4. Jason Yoo and Frank Wood. BayesPCN: A continually learnable predictive coding associative memory. In *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)*, 2022
5. Atılım Güneş Baydin, Barak A. Pearlmutter, Don Syme, Frank Wood, and Philip Torr. Gradients without backpropagation. *arXiv preprint arXiv:2202.08587*, 2022
6. Michael Teng, Michiel van de Panne, and Frank Wood. Exploration with multi-sample target values for distributional reinforcement learning. *arXiv preprint arXiv:2202.02693*, 2022
7. Peyman Bateni, Jarred Barber, Raghav Goyal, Vaden Masrani, Jan-Willem van de Meent, Leonid Sigal, and Frank Wood. Beyond simple meta-learning: Multi-purpose models for multi-domain, active and continual few-shot learning. *arXiv preprint arXiv:2201.05151*, 2022
8. Avi Pfeffer, Catherine Call, Frank Wood, Brad Rosenberg, Kirstin Bibbiani, Leonid Sigal, Ishaan Shah, Deniz Erdogmus, Sameer Singh, and Jan W van de Meent. Probabilistic label-efficient deep generative structures (pledges). Technical report, Charles River Analytics Inc., 2021
9. William Harvey, Saeid Naderiparizi, and Frank Wood. Image completion via inference in deep generative models. *arXiv preprint arXiv:2102.12037*, 2021
10. Rob Brekelmans, Vaden Masrani, Thang Bui, Frank Wood, Aram Galstyan, Greg Ver Steeg, and Frank Nielsen. Annealed importance sampling with q-paths. *arXiv preprint arXiv:2012.07823*, 2020
11. Andrew Warrington, J Wilder Lavington, Adam Ścibior, Mark Schmidt, and Frank Wood. Robust asymmetric learning in pomdps. *arXiv preprint arXiv:2012.15566*, 2020
12. William Harvey, Saeid Naderiparizi, and Frank Wood. Image completion via inference in deep generative models. *arXiv preprint arXiv:2102.12037*, 2021
13. Jason Yoo, Tony Joseph, Dylan Yung, S. Ali Nasser, and Frank Wood. Ensemble squared: A meta automl system. *arXiv preprint arXiv:2012.05390*, 2020
14. Saeid Naderiparizi, Kenny Chiu, Benjamin Bloem-Reddy, and Frank Wood. Uncertainty in neural processes. *arXiv preprint arXiv:2010.03753*, 2020
15. William Harvey, Michael Teng, and Frank Wood. Near-optimal glimpse sequences for improved hard attention neural network training. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2022
16. Peyman Bateni, Jarred Barber, Jan-Willem van de Meent, and Frank Wood. Improving few-shot visual classification with unlabelled examples. *arXiv preprint*, art. arXiv:2006.12245, 2020



17. Frank Wood, Michael Teng, and Rob Zinkov. Hasty-a generative model compiler. Technical report, University of Oxford Oxford United Kingdom, 2019
18. S Naderiparizi, A Ścibior, A Munk, M Ghadiri, AG Baydin, B Gram-Hansen, C Schroeder de Witt, R Zinkov, PHS Torr, T Rainforth, Y Whye Teh, and F Wood. Amortized rejection sampling in universal probabilistic programming. 2019. *arXiv preprint arXiv:1910.09056*
19. Michael Teng, Tuan Anh Le, Adam Ścibior, and Frank Wood. Imitation learning of factored multi-agent reactive models. *arXiv preprint arXiv:1903.04714*, 2019
20. 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
21. 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
22. 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
23. Michael Teng and Frank Wood. High throughput synchronous distributed stochastic gradient descent. *arXiv preprint arXiv:1803.04209*, 2018
24. Robert Cornish, Hongseok Yang, and Frank Wood. Towards a testable notion of generalization for generative adversarial networks. 2018
25. 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
26. Andrew Warrington and Frank Wood. Updating the vesicle-cnn synapse detector. *arXiv preprint arXiv:1710.11397*, 2017
27. 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
28. T. Rainforth, R. Cornish, H. Yang, and F. Wood. On the pitfalls of nested Monte Carlo. *arXiv preprint arXiv:1612.00951*, 2016
29. 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
30. T. Rainforth and F. Wood. Canonical correlation forests. *arXiv preprint arXiv:1507.05444*, 2015
31. Y. Perov and F. Wood. Learning probabilistic programs. *arXiv preprint arXiv:1407.2646*, 2014
32. Jonathan H Huggins and F. Wood. Infinite structured hidden semi-Markov models. *arXiv preprint arXiv:1407.0044*, 2014
33. J. W. van de Meent, B. Paige, and F. Wood. Tempering by subsampling. *arXiv preprint arXiv:1401.7145*, 2014
34. J. Zech and F. Wood. Inferring team strengths using a discrete Markov random field. *arXiv preprint arXiv:1305.1998*, 2013

35. W. Neiswanger and F. Wood. Unsupervised detection and tracking of arbitrary objects with dependent dirichlet process mixtures. *arXiv preprint arXiv:1210.3288*, 2012
36. 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. David Poole and Frank Wood. Probabilistic Programming Languages: Independent Choices and Deterministic Systems. Probabilistic and Causal Inference: The Works of Judea Pearl (1st ed.). Association for Computing Machinery, New York, NY, USA, 691–712. 2022.
2. 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**<sup>7</sup> 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**<sup>8</sup> 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.
- **Automated Pandemic Response Profiler**<sup>9</sup> is a web-based graphical user interface on top of a PyProb/FRED integration that enables automated exploration of policy interventions in response to COVID-19.

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

---

<sup>7</sup><https://anglican.ml>

<sup>8</sup><https://github.com/probprog/pyprob>

<sup>9</sup><https://covid19ideas.cs.ubc.ca/#/>