

Eric Nalisnick

Assistant Professor
University of Amsterdam
Amsterdam, Netherlands

e.t.nalisnick@uva.nl
enalisnick.github.io

Education

Ph.D. Computer Science	University of California, Irvine	2013-2018
M.S. Computer Science	Lehigh University	2012-2013
B.S. Computer Science & English Literature	Lehigh University	2008-2012

Employment

Assistant Professor, Tenure Track Supervisor: Max Welling	University of Amsterdam September 2020 to Present
Postdoctoral Research Associate Supervisor: José Miguel Hernández-Lobato	University of Cambridge September 2018 to September 2020
Research Scientist (one day per week)	DeepMind February 2019 to January 2020
Research Scientist Intern Supervisor: Balaji Lakshminarayanan	DeepMind Summer 2018
Graduate Student Researcher Supervisor: Padhraic Smyth	University of California, Irvine Fall 2013 to Spring 2018
Applied Scientist Intern Supervisors: Vijai Mohan, Eiman Elnahrawy	Amazon Fall 2016
Research Intern Supervisor: Hugo Larochelle	Twitter Summer 2016
Research Intern Supervisors: Rich Caruana, Nick Craswell	Microsoft Summer 2015
Research Scientist Intern Supervisors: Vijai Mohan, Rahul Bhagat	Amazon Summer 2014

Publications

* Denotes equal contribution

JOURNAL ARTICLES

1. G. Papamakarios*, **E. Nalisnick***, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing Flows for Probabilistic Modeling and Inference. *Journal of Machine Learning Research (JMLR)*, To Appear.

CONFERENCE ARTICLES

2. R. Pinsler, J. Gordon, **E. Nalisnick**, and J. M. Hernández-Lobato. Bayesian Batch Active Learning as Sparse Subset Approximation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
3. **E. Nalisnick**, J. M. Hernández-Lobato, and P. Smyth. Dropout as a Structured Shrinkage Prior. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
4. **E. Nalisnick***, A. Matsukawa*, Y. W. Teh, D. Gorur, and B. Lakshminarayanan. Hybrid Models with Deep and Invertible Features. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
5. **E. Nalisnick**, A. Matsukawa, Y. W. Teh, D. Gorur, and B. Lakshminarayanan. Do Deep Generative Models Know What They Don't Know? In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, 2019.
6. D. Ji, **E. Nalisnick**, Y. Qian, R. Scheuermann, and P. Smyth. Bayesian Trees for Automated Cytometry Data Analysis. In *Proceedings of Machine Learning for Healthcare (MLHC)*, 2018.
7. **E. Nalisnick** and P. Smyth. Learning Priors for Invariance. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
8. **E. Nalisnick** and P. Smyth. Learning Approximately Objective Priors. In *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2017.
9. **E. Nalisnick** and P. Smyth. Stick-Breaking Variational Autoencoders. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.
10. **E. Nalisnick**, B. Mitra, N. Craswell, and R. Caruana. Improving Document Ranking with Dual Word Embeddings. In *Proceedings of the 25th World Wide Web Conference (WWW)*, 2016.
11. **E. Nalisnick** and H. Baird. Character-to-Character Sentiment Analysis in Shakespeare's Plays. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*, 2013.
12. **E. Nalisnick** and H. Baird. Extracting Sentiment Networks from Shakespeare's Plays. In *Proceedings of the 12th International Conference on Document Analysis and Recognition (ICDAR)*, 2013.

PEER-REVIEWED WORKSHOP ARTICLES

13. Y. Zhang and **E. Nalisnick**. On the Inconsistency of Bayesian Inference for Misspecified Neural Networks. *Symposium on Advances in Approximate Bayesian Inference*, 2021.

14. E. Daxberger, **E. Nalisnick***, J. U. Allingham*, J. Antoran*, and J. M. Hernández-Lobato. Expressive yet Tractable Bayesian Deep Learning via Subnetwork Inference. *Symposium on Advances in Approximate Bayesian Inference*, 2021.
15. **E. Nalisnick**, J. Gordon, and J. M. Hernández-Lobato. Predictive Complexity Priors. *Uncertainty & Robustness in Deep Learning*, ICML 2020.
16. **E. Nalisnick**, A. Matsukawa, Y. W. Teh, and B. Lakshminarayanan. Detecting Out-of-Distribution Inputs to Deep Generative Models Using Typicality. *Bayesian Deep Learning*, NeurIPS 2019.
17. **E. Nalisnick** and J. M. Hernández-Lobato. Automatic Depth Determination for Bayesian ResNets. *Bayesian Deep Learning*, NeurIPS 2018.
18. **E. Nalisnick**, A. Matsukawa, Y.W. Teh, D. Gorur, and B. Lakshminarayanan. Do Deep Generative Models Know What They Don't Know? *Bayesian Deep Learning*, NeurIPS 2018.
19. **E. Nalisnick***, A. Matsukawa*, Y.W. Teh, D. Gorur, and B. Lakshminarayanan. Hybrid Models with Deep and Invertible Features. *Bayesian Deep Learning*, NeurIPS 2018.
20. O. Rybakov, V. Mohan, A. Misra, S. LeGrand, R. Joseph, K. Chung, S. Singh, Q. You, **E. Nalisnick**, L. Dirac, and R. Luo. The Effectiveness of a Two-Layer Neural Network for Recommendations. Workshop Track, ICLR 2018.
21. D. Ji, **E. Nalisnick**, and P. Smyth. Mondrian Processes for Flow Cytometry Analysis. *Machine Learning for Health*, NeurIPS 2017.
22. **E. Nalisnick** and P. Smyth. Variational Inference with Stein Mixtures. *Advances in Approximate Bayesian Inference*, NIPS 2017.
23. **E. Nalisnick** and P. Smyth. The Amortized Bootstrap. *Implicit Models*, ICML 2017.
24. **E. Nalisnick** and P. Smyth. Variational Reference Priors. Workshop Track, ICLR 2017.
25. **E. Nalisnick**, L. Hertel, and P. Smyth. Approximate Inference for Deep Latent Gaussian Mixtures. *Bayesian Deep Learning*, NeurIPS 2016.
26. **E. Nalisnick** and P. Smyth. Nonparametric Deep Generative Models with Stick-Breaking Priors. *Data-Efficient Machine Learning*, ICML 2016.
27. J. Park, M. Blume-Kohout, R. Krestel, **E. Nalisnick**, and P. Smyth. Analyzing NIH Funding Patterns over Time with Statistical Text Analysis. *Scholarly Big Data*, AAAI 2016.

THESES

1. **E. Nalisnick**. On Priors for Bayesian Neural Networks. *Doctoral Dissertation*, University of California, Irvine, 2018.
2. **E. Nalisnick**. Automatic Methods for Tracking Sentiment Dynamics in Plays. *Master's Thesis*, Lehigh University, 2013.
3. **E. Nalisnick**. A Combinatorial Explanation for a Conjecture of Fomin and Zelevinsky. *Honors Thesis*, Lehigh University, 2012.

Teaching

1. *Learning* (“Leren”) | University of Amsterdam 2020
Introduction to Machine Learning, 181 Undergraduate Students, 15 Teaching Assistants

Awards

1. Top Reviewer NeurIPS 2017, ICML 2019, ICML 2020
2. Travel Award ACL 2013, ICML Workshop on Implicit Models 2017, AISTats 2018, ICML 2019
3. NSF Graduate Research Fellowship — Honorable Mention 2014
4. UCI Graduate Dean’s Recruitment Fellowship 2013
5. Phi Beta Kappa 2012

Professional Service

WORKSHOP ORGANIZATION

- | | |
|---|------|
| Bayesian Deep Learning, ELLIS Workshop / NeurIPS Meetup | 2020 |
| Bayesian Deep Learning, NeurIPS Workshop | 2019 |

AREA CHAIR / SENIOR PROGRAM COMMITTEE

- | | |
|---|------|
| International Conference on Learning Representations (ICLR) | 2021 |
|---|------|

JOURNAL REVIEWING

- | | |
|-------------------------------------|-----------------|
| Machine Learning Research | 2018 to Present |
| Neural Processing Letters | 2019 |
| Machine Learning | 2017 |
| Data Mining and Knowledge Discovery | 2017 |

CONFERENCE REVIEWING

- | | |
|---|-----------------|
| Neural Information Processing Systems (NeurIPS) | 2016 to Present |
| International Conference on Learning Representations (ICLR) | 2018 to Present |
| International Conference on Machine Learning (ICML) | 2018 to Present |
| Artificial Intelligence and Statistics (AISTats) | 2019 to Present |
| Uncertainty in Artificial Intelligence (UAI) | 2019 to Present |
| Association for the Advancement of Artificial Intelligence (AAAI) | 2020 to Present |
| International Joint Conference on Artificial Intelligence (IJCAI) | 2019 |

WORKSHOP REVIEWING

Advances in Approximate Bayesian Inference	2018 to Present
I Cant Believe It's Not Better!	2020
Uncertainty & Robustness in Deep Learning	2020
Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models	2020

Invited Talks

1. Detecting Distribution Shift with Deep Generative Models, SYDNEY ML MEETUP 2020
2. Detecting Distribution Shift with Deep Generative Models, INNFP+, ICML WORKSHOP 2020
3. Building and Critiquing Models for Probabilistic Deep Learning, GATSBY UNIT, UCL 2020
4. Building and Critiquing Models for Probabilistic Deep Learning, CARNEGIE MELLON UNIV. 2020
5. Building and Critiquing Models for Probabilistic Deep Learning, UNIV. OF NORTH CAROLINA 2020
6. Deep Learning & Statistics: Bridging the Gap with Prob. Structure, UNIV. OF AMSTERDAM 2020
7. Deep Learning & Statistics: Bridging the Gap with Prob. Structure, UC SANTA BARBARA 2020
8. Deep Learning Under Covariate Shift, UCI AI/ML SEMINAR 2019
9. Normalizing Flows for Tractable Probabilistic Modeling and Inference, T-PRIME, NEURIPS 2019
10. Deep Learning: A Synthesis from Probabilistic Foundations, RAND CORP STATS. SEMINAR 2019
11. Evaluating Deep Generative Models on Out-of-Distribution Inputs, OXFORD STATS. SEMINAR 2019
12. Do Deep Generative Models Know What They Don't Know?, CAMAIML (MSR CAMBRIDGE) 2019
13. Do Deep Generative Models Know What They Don't Know?, CAMBRIDGE LTL SEMINAR 2019
14. Structured Shrinkage Priors for Neural Networks, IMPERIAL COLLEGE STATISTICS SEMINAR 2018
15. Deep Learning: A Synthesis from Probabilistic Foundations, UCI STATISTICS SEMINAR 2018
16. Approximate Inference for Frequentist Uncertainty Estimation, SOCAL ML SYMPOSIUM 2017
17. The Amortized Bootstrap, ICML WORKSHOP ON IMPLICIT MODELS 2017
18. Deep Generative Models with Stick-Breaking Priors, UCI AI/ML SEMINAR 2017
19. Alternative Priors for Deep Generative Models, OPENAI 2017
20. Nonparametric Deep Generative Models, ICML WORKSHOP ON DATA-EFFICIENT ML 2016