Eric Nalisnick

nalisnick@jhu.edu

enalisnick.github.io

Johns Hopkins University

University of Amsterdam

University of Cambridge

March 2024 - present

Academic Employment

Assistant Professor, Tenure Track Department of Computer Science

Assistant Professor, Tenure Track Amsterdam Machine Learning Lab

Postdoctoral Research Associate Computational and Biological Learning Laboratory

Industrial Employment

Research Scientist Part time, one day per week (20% FTE)

Research Scientist Intern Supervisor: Balaji Lakshminarayanan

Applied Scientist Intern Supervisors: Vijai Mohan, Eiman Elnahrawy

Research Intern Supervisor: Hugo Larochelle

Research Intern Supervisors: Rich Caruana, Nick Craswell

Research Scientist Intern Supervisors: Vijai Mohan, Rahul Bhagat

Education

Ph.D. Computer Science	2013 - 2018
University of California, Irvine	
Dissertation: On Priors for Bayesian Neural Networks	
Advisor: Padhraic Smyth	
M.S. Computer Science	2012 - 2013
Lehigh University	
Thesis: Automatic Methods for Tracking Sentiment Dynamics in Plays	
Advisor: Henry S. Baird	

September 2018 to September 2020

September 2020 - February 2024

Google DeepMind February 2019 to January 2020

> Google DeepMind Summer 2018

> > Amazon Fall 2016

Twitter Summer 2016

Microsoft Summer 2015

Amazon Summer 2014

B.S. Computer Science & English Literature **Lehigh University** Honors Thesis: *A Combinatorial Explanation for a Conjecture of Fomin & Zelevinsky* Advisor: Mark Skandera

Academic Honors and Awards

Google Award for Inclusion Research	2024
Notable paper (32 of 496 accepted papers (6%))	AISTATS 2023
Veni Laureate, Dutch Research Council (NWO)	2021
ELLIS Scholar, European Lab for Learning and Intelligent System	s Society 2021
Top / Best Reviewer	NeurIPS 2017, ICML 2019, ICML 2020
Honorable Mention, Best Poster Award	SoCal ML Symposium 2017

Academic Supervision

CURRENT PHD CANDIDATES

Rajeev Verma	University of Amsterdam, 2023 -
Alexander Timans	University of Amsterdam, 2022 -
Metod Jazbec	University of Amsterdam, 2022 -
Mona Schirmer	University of Amsterdam, 2022 -
Dharmesh Tailor	University of Amsterdam, 2021 -
Saba Amiri (with Adam Belloum and Sander Klous)	University of Amsterdam, 2021 -
Putra Manggala (with Holger Hoos)	University of Amsterdam, 2021 -
Urja Khurana (with Antske Fokkens)	Vrije University Amsterdam, 2020 -
James Allingham (with José M. Hernández-Lobato via ELLIS)	University of Cambridge, 2019 - 2024
Former PhD Candidates	
Mrinank Sharma (with Tom Rainforth and Yee Whye Teh)	University of Oxford

Thesis: Uncertainty Estimation with Small and Large Models First position after graduation: Research Scientist, Anthropic 2020 - 2024

MASTERS STUDENTS

University of Amsterdam: Jaap Stefels (2023), Thomas Jurriaans (2023), Mark Fokkema (2023), Arsen Sheverdin (2023), Nils Lehmann (2022), Shuai Wang (2022), Rajeev Verma (2022), Daniël Nobbe (2021).

VISITING STUDENTS

University of Amsterdam: Rutger Hendrix (2023), Javier Antorán (2022), Daniel Barrejón (2022).

2008 - 2012

Eric Nalisnick

Teaching

Human-in-the-Loop Machine Learning (Graduate)	Johns Hopkins University, 2024
Human-in-the-Loop Machine Learning (Graduate)	University of Amsterdam, 2023
Machine Learning I (Graduate)	University of Amsterdam, 2023
Introduction to Machine Learning (Undergraduate)	University of Amsterdam, 2020 - 2022
Bayesian Deep Learning Module, Deep Learning II (Graduate)	University of Amsterdam, 2022 - 2023
Project in AI (Graduate)	University of Amsterdam, 2021

Professional Service

Organization / Leadership Roles	
Advisory Committee on Data Science and AI, American Statistical Association	2024 - present
Tractable Probabilistic Modeling, UAI Workshop	2022 - 2023
Anomaly Detection for Scientific Discovery, Virtual Seminar Series	2021 - 2022
Bayesian Deep Learning, NeurIPS Workshop	2019, 2021
Bayesian Deep Learning, ELLIS Workshop / NeurIPS Meetup	2020
Area Chair / Senior Program Committee	
Neural Information Processing Systems (NeurIPS)	2021 - present
International Conference on Machine Learning (ICML)	2021 - present
Uncertainty in Artificial Intelligence (UAI)	2021 - present
Artificial Intelligence and Statistics (AIStats)	2022
International Conference on Learning Representations (ICLR)	2021
Journal Reviewing	
American Statistical Association (JASA)	2020 - present
Machine Learning Research (JMLR)	2018 - present
Journal of Econometrics	2022
Advances in Statistical Analysis	2020
Neural Processing Letters	2019
Conference Reviewing	
Conference on Computer Vision and Pattern Recognition (CVPR)	2024
Neural Information Processing Systems (NeurIPS)	2016 - 2021

International Conference on Machine Learning (ICML) Artificial Intelligence and Statistics (AIStats) Uncertainty in Artificial Intelligence (UAI) Hybrid Human-Artificial Intelligence (HHAI) Association for the Advancement of Artificial Intelligence (AAAI) International Joint Conference on Artificial Intelligence (IJCAI) GRANT REVIEWING Johns Hopkins Discovery Awards U.SIsrael Binational Science Foundation Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Baversian Outimation on Non-Conventional Sparch Sparces	2018 - 2021 2019 - 2021 2019 - 2020
Artificial Intelligence and Statistics (AIStats) Uncertainty in Artificial Intelligence (UAI) Hybrid Human-Artificial Intelligence (HHAI) Association for the Advancement of Artificial Intelligence (AAAI) International Joint Conference on Artificial Intelligence (IJCAI) GRANT REVIEWING Johns Hopkins Discovery Awards U.SIsrael Binational Science Foundation Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam	2019 - 2021
Uncertainty in Artificial Intelligence (UAI) Hybrid Human-Artificial Intelligence (HHAI) Association for the Advancement of Artificial Intelligence (AAAI) International Joint Conference on Artificial Intelligence (IJCAI) GRANT REVIEWING Johns Hopkins Discovery Awards U.SIsrael Binational Science Foundation Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bavesian Ontimization on Non-Conventional Search Spaces	2010 - 2020
Hybrid Human-Artificial Intelligence (HHAI) Association for the Advancement of Artificial Intelligence (AAAI) International Joint Conference on Artificial Intelligence (IJCAI) GRANT REVIEWING Johns Hopkins Discovery Awards U.SIsrael Binational Science Foundation Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Bertrand Charpentier, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Baversion Ohtimization on Non-Conventional Search Spaces	2017 - 2020
Association for the Advancement of Artificial Intelligence (AAAI) International Joint Conference on Artificial Intelligence (IJCAI) GRANT REVIEWING Johns Hopkins Discovery Awards U.SIsrael Binational Science Foundation Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bavesian Ontimization on Non-Conventional Search Spaces	2022
International Joint Conference on Artificial Intelligence (IJCAI) GRANT REVIEWING Johns Hopkins Discovery Awards U.SIsrael Binational Science Foundation Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bavesian Ontimization on Non-Conventional Search Spaces	2020 - 2021
GRANT REVIEWING Johns Hopkins Discovery Awards U.SIsrael Binational Science Foundation Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bavesian Ontimization on Non-Conventional Search Spaces	2019
Johns Hopkins Discovery Awards U.SIsrael Binational Science Foundation Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bavesian Optimization on Nan-Conventional Search Spaces	
U.SIsrael Binational Science Foundation Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bayesian Ontimization on Non-Conventional Search Spaces	2024
 Institutional Service Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Baversian Optimization on Non-Conventioned Search Spaces 	2023
Computer science curriculum committee, Johns Hopkins University AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bayesian Optimization on Non-Conventional Search Spaces	
 AI education program committee, University of Amsterdam Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark <i>Uncertainty and the Medical Interview</i> Bertrand Charpentier, Technical University of Munich <i>Uncertainty Estimation for Independent and Non-Independent Data</i> Kamil Deja, Warsaw University of Technology <i>Data Representations in Generative Modelling</i> ChangYong Oh, University of Amsterdam Bavesian Ontimization on Non-Conventional Search Spaces 	2024 - present
 Faculty hiring committee, Machine Learning Lab, University of Amsterdam Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bayesian Optimization on Non-Conventional Search Spaces 	2021 - 2024
 Doctoral Committees Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bayesian Optimization on Non-Conventional Search Spaces 	2021
Jakob Havtorn, Technical University of Denmark Uncertainty and the Medical Interview Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bayesian Optimization on Non-Conventional Search Spaces	
Bertrand Charpentier, Technical University of Munich Uncertainty Estimation for Independent and Non-Independent Data Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bayesian Optimization on Non-Conventional Search Spaces	2024
Kamil Deja, Warsaw University of Technology Data Representations in Generative Modelling ChangYong Oh, University of Amsterdam Bayesian Optimization on Non-Conventional Search Spaces	2024
ChangYong Oh, University of Amsterdam Bayesian Optimization on Non-Conventional Search Spaces	2023
Dayesian Optimization on 110h Conventional Search Spaces	
Emiel Hoogeboom, University of Amsterdam Normalizing Flows and Diffusion Models for Discrete and Geometric Data	2023
Shi Hu, University of Amsterdam Uncertainty, Robustness and Safety in Artificial Intelligence	2023 2023

Publications

* Denotes equal contribution

JOURNAL ARTICLES

- 1. E. Nalisnick, P. Smyth, and D. Tran. A Brief Tour of Deep Learning from a Statistical Perspective. *Annual Review of Statistics and Its Application*, 2023.
- 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), 2021.

CONFERENCE ARTICLES

- 3. M. Jazbec, P. Forré, S. Mandt, D. Zhang, and **E. Nalisnick**. Anytime-Valid Confidence Sequences for Consistent Uncertainty Estimation in Early-Exit Neural Networks. In *Proceedings of the 40th Conference on Uncertainty in Artificial Intelligence* (UAI), 2024.
- D. Tailor, A. Patra, R. Verma, P. Manggala, and E. Nalisnick. Learning to Defer to a Population: A Meta-Learning Approach. In Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS), 2024. Oral presentation.
- M. Jazbec, J. U. Allingham, D. Zhang, and E. Nalisnick. Towards Anytime Classification in Early-Exit Architectures by Enforcing Conditional Monotonicity. In Advances in Neural Information Processing Systems (NeurIPS), 2023.
- 6. D. Tailor, M. E. Khan, and **E. Nalisnick**. Exploiting Inferential Structure in Neural Processes. In *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence* (UAI), 2023.
- J. Antoran*, S. Padhy*, R. Barbano, E. Nalisnick, D. Janz, and J. M. Hernández-Lobato. Sampling-Based Inference for Large Linear Models, with Application to Linearised Laplace. In Proceedings of the 11th International Conference on Learning Representations (ICLR), 2023.
- R. Verma*, D. Barrejón*, and E. Nalisnick. Learning to Defer to Multiple Experts: Consistent Surrogate Losses, Confidence Calibration, and Conformal Ensembles. In Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS), 2023.
- 9. M. Sharma, S. Farquhar, E. Nalisnick, and T. Rainforth. Do Bayesian Neural Networks Need To Be Fully Stochastic? In Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS), 2023. 'Notable' oral presentation: 6% of accepted papers.
- 10. R. Verma and E. Nalisnick. Calibrated Learning to Defer with One-vs-All Classifiers. In Proceedings of the 39th International Conference on Machine Learning (ICML), 2022.
- J. Antoran, D. Janz, J. U. Allingham, E. Daxberger, R. Barbano, E. Nalisnick, and J. M. Hernández-Lobato. Adapting the Linearised Laplace Model Evidence for Modern Deep Learning. In Proceedings of the 39th International Conference on Machine Learning (ICML), 2022.
- 12. E. Daxberger, E. Nalisnick*, J. U. Allingham*, J. Antoran*, and J. M. Hernández-Lobato. Expressive yet Tractable Bayesian Deep Learning via Subnetwork Inference. In *Proceedings of the 38th International Conference on Machine Learning* (ICML), 2021.

- 13. E. Nalisnick, J. Gordon, and J. M. Hernández-Lobato. Predictive Complexity Priors. In Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS), 2021.
- 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.
- 15. 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.
- 16. 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.
- 17. 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.
- 18. 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.
- 19. E. Nalisnick and P. Smyth. Learning Priors for Invariance. In Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS), 2018.
- 20. E. Nalisnick and P. Smyth. Learning Approximately Objective Priors. In Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence (UAI), 2017.
- 21. E. Nalisnick and P. Smyth. Stick-Breaking Variational Autoencoders. In Proceedings of the 5th International Conference on Learning Representations (ICLR), 2017.
- 22. 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.
- 23. 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.
- 24. 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 (WITHOUT CONFERENCE VERSION)

- N. Lehmann, N. M. Gottschling, S. Depeweg, and E. Nalisnick. Uncertainty Aware Tropical Cyclone Wind Speed Estimation from Satellite Data. *Machine Learning for Remote Sensing*, ICLR 2024. *Oral* presentation.
- 26. M. Schirmer, D. Zhang, and **E. Nalisnick**. Beyond Top-Class Agreement: Using Divergences to Forecast Performance under Distribution Shift. *Distribution Shifts*, NeurIPS 2023.
- 27. A. Timans, C.-N. Straehle, K. Sakmann, and **E. Nalisnick**. Adaptive Bounding Box Uncertainty via Conformal Prediction. *Uncertainty Quantification for Computer Vision*, ICCV 2023.

- 28. N. Lehmann, N. M. Gottschling, S. Depeweg, and **E. Nalisnick**. A Comparison of Uncertainty Quantification Methods for Earth Observation Image Regression Data. *Uncertainty Quantification for Computer Vision*, ICCV 2023.
- 29. J. U. Allingham, J. Antoran, S. Padhy, E. Nalisnick, and J. M. Hernández-Lobato. Learning Generative Models with Invariance to Symmetries. *Symmetry and Geometry in Neural Representations*, NeurIPS 2022.
- 30. S. Wang and E. Nalisnick. Active Learning for Multilingual Fingerspelling Corpora. *Adaptive Experimental Design and Active Learning in the Real World*, ICML 2022.
- 31. J. U. Allingham and E. Nalisnick. A Product of Experts Approach to Early-Exit Ensembles. *Dynamic Neural Networks*, ICML 2022.
- 32. P. Manggala, H. Hoos, and E. Nalisnick. Bayesian Weak Supervision via an Optimal Transport Approach. *Human-Machine Collaboration and Teaming*, ICML 2022.
- 33. U. Khurana, I. Vermeulen, **E. Nalisnick**, M. van Noorloos, and A. Fokkens. Hate Speech Criteria: A Modular Approach to Task-Specific Hate Speech Definitions. *Online Abuse and Harms*, NAACL 2022.
- 34. S. Amiri, A. Belloum, **E. Nalisnick**, S. Klous, and L. Gommans. On the Impact of Non-IID Data on the Performance and Fairness of Differentially Private Federated Learning. *Dependable and Secure Machine Learning*, DSN 2022.
- 35. U. Khurana, **E. Nalisnick**, and A. Fokkens. How Emotionally Stable is ALBERT? Testing Robustness with Stochastic Weight Averaging on a Sentiment Analysis Task. *Evaluation and Comparison of NLP Systems*, EMNLP 2021.
- 36. P. Manggala, H. Hoos, and **E. Nalisnick**. Bayesian Regression from Multiple Sources of Weak Supervision. *Machine Learning for Data*, ICML 2021.
- 37. Y. Zhang and **E. Nalisnick**. On the Inconsistency of Bayesian Inference for Misspecified Neural Networks. *Symposium on Advances in Approximate Bayesian Inference*, 2021.
- 38. 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.
- 39. E. Nalisnick and J. M. Hernández-Lobato. Automatic Depth Determination for Bayesian ResNets. *Bayesian Deep Learning*, NeurIPS 2018.
- 40. 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.
- 41. D. Ji, E. Nalisnick, and P. Smyth. Mondrian Processes for Flow Cytometry Analysis. *Machine Learning for Health*, NeurIPS 2017.
- 42. E. Nalisnick and P. Smyth. Variational Inference with Stein Mixtures. *Advances in Approximate Bayesian Inference*, NIPS 2017.
- 43. E. Nalisnick and P. Smyth. The Amortized Bootstrap. Implicit Models, ICML 2017. Oral presentation.

- 44. E. Nalisnick, L. Hertel, and P. Smyth. Approximate Inference for Deep Latent Gaussian Mixtures. *Bayesian Deep Learning*, NeurIPS 2016.
- 45. 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.

Patents

1. E. M. H. Elnahrawy, V. Mohan, and **E. Nalisnick**. Generation and Use of Model Parameters in Cold-Start Scenarios. U.S. Patent Number 10,726,334. 28 July 2020.

Invited Talks

INTERNATIONAL VENUES, LARGE NATIONAL VENUES, KEYNOTES

1.	Detecting Distribution Shift with Deep Generative Models Dutch Society of Pattern Recognition and Image Processing	2023 Amsterdam, Netherlands
2.	Towards Anytime Uncertainty Estimation in Early-Exit Neural Networks Workshop on <i>Uncertainty Quantification for Computer Vision</i>	2023 Dorio Eronoo
0	INTERNATIONAL CONFERENCE ON COMPUTER VISION (ICCV)	Paris, France
3.	ELLIS Workshop on Robust Machine Learning (RobustML)	2023 Helsinki, Finland
4.	Bayesian Learning & Uncertainty Quantification Mediterranean Machine Learning Summer School	2023 Thessaloniki, Greece
5.	Towards Anytime Computation in Deep Architectures Dagstuhl Seminar on Deep Generative Models	2023 Wadern, Germany
6.	Towards Informative Priors for Bayesian Deep Learning Dagstuhl Seminar on Tractable Probabilistic Models	2022 Wadern, Germany
7.	Predictive Complexity Priors GenU Workshop: Gen. Models and Uncertainty Quantification	2021 Copenhagen, Denmark
8.	Detecting Distribution Shift with Deep Generative Models Workshop on <i>Invertible Neural Nets, Normalizing Flows, and Explicit Likele</i> INTERNATIONAL CONFERENCE ON MACHINE LEARNING (ICML)	2020 ihood Models (INNF+) Virtual
9.	Normalizing Flows for Tractable Probabilistic Modeling and Inference T-PRIME, NEURIPS SOCIAL	2019 Vancouver, Canada
Depa	rtment Seminars, Companies, Local Events	
1.	AI Risks: From Today to Doomsday — An Academic Panel Discussion Amsterdam ELLIS Unit	2023
2.	Learning to Defer to One, Multiple, or a Population of Expert(s) Amazon	2023

3.	Towards a Statistical Foundation for Human-AI Collaboration	2023
	Univ. of Tübingen / Bosch AI, Oregon State Univ.: Computer Science, Johns Hopkins Computer Science, George Washington Univ.: Computer Science	Univ.:
4.	On the Calibration of Learning-to-Defer Systems Univ. of California, Riverside: Computer Science, Univ. of Manchester: Statistics	2022
5.	Predictive Complexity Priors Univ. of Edinburgh: ANC Seminar, Imperial College: Statistics, Alan Turing Institu	2021 ге
6.	Detecting Distribution Shift with Deep Generative Models Sydney Machine Learning Meetup	2020
7.	Building and Critiquing Models for Probabilistic Deep Learning UNIVERSITY COLLEGE LONDON: GATSBY UNIT, CARNEGIE MELLON UNIV.: STATISTICS, UNI NORTH CAROLINA: COMPUTER SCIENCE, UNIV. OF AMSTERDAM: INFORMATICS	2020 v. of
8.	Deep Learning & Statistics: Bridging the Gap with Prob. Structure UNIV. OF CALIFORNIA, SANTA BARBARA: STATISTICS	2020
9.	Deep Learning Under Covariate Shift Univ. of California, Irvine: AI/ML Seminar	2019
10.	Deep Learning: A Synthesis from Probabilistic Foundations2018RAND Corporation: Statistics, Univ. of California, Irvine: Statistics2018	- 2019
11.	Evaluating Deep Generative Models on Out-of-Distribution Inputs UNIV. OF OXFORD: STATISTICS, CAMAIML (MSR CAMBRIDGE), UNIV. OF CAMBRIDGE: LTL SEM	2019 inar
12.	Structured Shrinkage Priors for Neural Networks Imperial College: Statistics	2018
13.	Approximate Inference for Frequentist Uncertainty Estimation SoCAL ML SYMPOSIUM	2017
14.	Deep Generative Models with Stick-Breaking Priors UNIV. OF CALIFORNIA, IRVINE: AI/ML SEMINAR, OPENAI	2017