Eric Nalisnick

nalisnick@jhu.edu enalisnick.github.io

Academic Employment

Assistant Professor, Tenure Track — Johns Hopkins University
Department of Computer Science — March 2024 - present

Assistant Professor, Tenure Track — University of Amsterdam
Amsterdam Machine Learning Lab — September 2020 - February 2024

Postdoctoral Research Associate — University of Cambridge
Computational and Biological Learning Laboratory — September 2018 to September 2020

Industrial Employment

Research Scientist — Google DeepMind
Part time, one day per week (20% FTE) — February 2019 to January 2020

Research Scientist Intern — Google DeepMind
Supervisor: Balaji Lakshminarayanan — Summer 2018

Applied Scientist Intern — Amazon
Supervisors: Vijai Mohan, Eiman Elnahrawy — Fall 2016

Research Intern — Twitter
Supervisor: Hugo Larochelle — Summer 2016

Research Intern — Microsoft
Supervisors: Rich Caruana, Nick Craswell — Summer 2015

Research Scientist Intern — Amazon
Supervisors: Vijai Mohan, Rahul Bhagat — Summer 2014

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

B.S. Computer Science & English Literature 2008 - 2012
Lehigh University
Honors Thesis: *A Combinatorial Explanation for a Conjecture of Fomin & Zelevinsky*
Advisor: Mark Skandera

Honors and Awards

Joel Dean Excellence in Teaching Award, JHU Computer Science Dept. 2026

Best Paper (2 of 38 accepted papers), QUESTION Workshop at ICLR 2025

Oral presentation (8% of accepted papers), ECCV 2024

Outstanding Student Paper (as supervisor, 7 of 546 accepted papers), AISTATS 2024

Google Award for Inclusion Research 2024

Notable Paper (32 of 496 accepted papers), AISTATS 2023

Veni Laureate, Dutch Research Council 2021

Scholar, European Lab for Learning and Intelligent Systems (ELLIS) Society 2021

Top / Best Reviewer NeurIPS 2017, ICML 2019, ICML 2020

Honorable Mention for Best Poster Award, SoCal ML Symposium 2017

Academic Supervision

Current PhD Students / Candidates

Luisa Li Johns Hopkins University, 2026 -

Chi Zhang (with Kimia Ghobadi) Johns Hopkins University, 2025 -

Xi Wang Johns Hopkins University, 2025 -

Andrea Wynn (with Anqi Liu) Johns Hopkins University, 2024 -

Rajeev Verma (with Christian Naesseth, Volker Fischer) University of Amsterdam, 2023 -

Alexander Timans University of Amsterdam, 2022 -
(with Christian Naesseth, Kaspar Sakmann, Christoph-Nikolas Straehle)

Metod Jazbec (with Dan Zhang) University of Amsterdam, 2022 -

Mona Schirmer (with Dan Zhang) University of Amsterdam, 2022 -

Dharmesh Tailor University of Amsterdam, 2021 -

Putra Manggala (with Holger Hoos) University of Amsterdam, 2021 -

Current Doctor of Engineering Candidates

Matthew Renze Johns Hopkins University, 2025 -

Graduated PhD Candidates

Saba Amiri (with Adam Belloum and Sander Klous) University of Amsterdam
Thesis: *From Patient to Pattern: Probabilistic Models for Synthetic Data in Digital Health* 2021 - 2026
First position after graduation: Applied Scientist, KPMG

Urja Khurana (with Antske Fokkens) Vrije University Amsterdam
Thesis: *From Concept to Impact: Toward More Robust Language Model Deployment* 2020 - 2025
First position after graduation: Postdoctoral Researcher, Delft University of Technology (TU Delft)

James Allingham (external advisor via ELLIS) University of Cambridge
Thesis: *Improving Deep Learning with Probabilistic Approaches* 2019 - 2024
First position after graduation: Research Scientist, Google DeepMind

Mrinank Sharma (with Tom Rainforth and Yee Whye Teh) University of Oxford
Thesis: *Uncertainty Estimation with Small and Large Models* 2020 - 2024
First position after graduation: Research Scientist, Anthropic

Masters Students

Johns Hopkins University: Mirra Klimov (2026), Evie Kim (2026), Ziyao Mou (2025-2026), Rishi More (2025), Han Liu (2024-2026), Yizirui Fang (2024-2025), Runzhou Chen (2024), Kunlun Li (2024)

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).

Undergraduate Students

Johns Hopkins University: Arda Uzunoğlu (2026 - present), Tomoya Furutani (2026 - present), Alessa Carbo (2024 - present), Sophia Qian (2026), Angela Hu (2025)

Visiting Students

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

Teaching

Human-in-the-Loop Machine Learning (Graduate) Johns Hopkins University, 2024 - present

Machine Learning: Deep Learning (Graduate) Johns Hopkins University, 2025 - 2026

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
AI Safety: Towards Trustworthy Agentic AI, UAI Workshop	2026
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
Artificial Intelligence and Statistics (AIStats)	2022, 2027
Uncertainty in Artificial Intelligence (UAI)	2021 - 2025
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

International Conference on Learning Representations (ICLR)	2018 - 2024, 2027
Conference on Computer Vision and Pattern Recognition (CVPR)	2024
Neural Information Processing Systems (NeurIPS)	2016 - 2021
International Conference on Machine Learning (ICML)	2018 - 2021
Artificial Intelligence and Statistics (AIStats)	2019 - 2021
Uncertainty in Artificial Intelligence (UAI)	2019 - 2020
Hybrid Human-Artificial Intelligence (HHAI)	2022
Association for the Advancement of Artificial Intelligence (AAAI)	2020 - 2021
International Joint Conference on Artificial Intelligence (IJCAI)	2019

Grant Reviewing

Johns Hopkins Provost's Undergraduate Research Award	2026
European Research Council	2025
US Army Research Office	2024
Johns Hopkins Discovery Awards	2024
U.S.-Israel Binational Science Foundation	2023

Institutional Service

Data Science & AI faculty hiring committee, Johns Hopkins University	2025 - 2027
Summer internship program committee, JHU Data Science & AI Institute	2026 - present
Computer Science curriculum committee, Johns Hopkins University	2024 - present
Cognitive Science faculty hiring committee, Johns Hopkins University	2024
AI education program committee, University of Amsterdam	2021 - 2024
Faculty hiring committee, Machine Learning Lab, University of Amsterdam	2021

Doctoral Committees

Dissertation Defense Examinations:

Richard Williams, University of California, Los Angeles 2026
Efficient Generative Modeling of Complex Graph Data and Posterior Distributions over Phylogenetic Trees

Alain Ryser, ETH Zurich 2025
Incorporating Domain Knowledge in Self-supervised Rep. Learning for Clustering and Anomaly Detection

Teodora Pandeva, University of Amsterdam 2025
Machine Learning for Multi-Source Data Integration

Federico Bergamin, Technical University of Denmark 2024
Advances in Deep Generative Models, Approximate Inference, and their Applications

Jakob Havtorn, Technical University of Denmark 2024
Uncertainty and the Medical Interview

Bertrand Charpentier, Technical University of Munich 2024
Uncertainty Estimation for Independent and Non-Independent Data

Kamil Deja, Warsaw University of Technology 2023
Data Representations in Generative Modelling

ChangYong Oh, University of Amsterdam 2023
Bayesian Optimization on Non-Conventional Search Spaces

Emiel Hoogeboom, University of Amsterdam 2023
Normalizing Flows and Diffusion Models for Discrete and Geometric Data

Shi Hu, University of Amsterdam 2022
Uncertainty, Robustness and Safety in Artificial Intelligence

Doctoral Candidacy Examinations (not as advisor): Leonard Bereska (Univ. of Amsterdam, 2023), Ilze Amanda Auzina (Univ. of Amsterdam, 2023), Zhengping Jiang (JHU, 2025), Gina Wong (JHU, 2025), Aayush Mishra (JHU, 2025), Richard Williams (UCLA, 2025), Ha Manh Bui (JHU, 2026), Thomas Shaddix (JHU, 2026), Tianjian Li (JHU, 2026)

Publications

* Denotes equal contribution

Journal Articles

1. X. Wang and **E. Nalisnick**. Are Vision Language Models Robust to Classic Uncertainty Challenges? In *Transactions on Machine Learning Research* (TMLR), 2026.
2. Y. Fang and **E. Nalisnick**. Learning to Defer with an Uncertain Rejector via Conformal Prediction. In *Transactions on Machine Learning Research* (TMLR), 2026.
3. M. Schirmer, D. Zhang, and **E. Nalisnick**. Temporal Test-Time Adaptation with State-Space Models. In *Transactions on Machine Learning Research* (TMLR), 2025. *Journal-to-Conference (J2C) Certification (∼10% of accepted papers).*[1]
4. R. Williams, **E. Nalisnick**, and A. Holbrook. Scalable Generative Modeling of Weighted Graphs. In *Transactions on Machine Learning Research* (TMLR), 2025.
5. L. Manduchi, K. Pandey, C. Meister, R. Bamler, R. Cotterell, S. Däubener, S. Fellenz, A. Fischer, T. Gärtner, M. Kirchler, M. Kloft, Y. Li, C. Lippert, G. de Melo, **E. Nalisnick**, B. Ommer, R. Ranganath, M. Rudolph, K. Ullrich, G. Van den Broeck, J. E. Vogt, Y. Wang, F. Wenzel, F. Wood, S. Mandt, and V. Fortuin. On the Challenges and Opportunities in Generative AI. In *Transactions on Machine Learning Research* (TMLR), 2025.
6. N. Lehmann, N. M. Gottschling, J. Gawlikowski, A. Stewart, S. Depeweg, and **E. Nalisnick**. Lightning UQ Box: Uncertainty Quantification for Neural Networks. *Journal of Machine Learning Research* (JMLR), Machine Learning Open Source Software track, 2025.
7. S. Amiri, **E. Nalisnick**, A. Belloum, S. Klous, L. Gommans. Practical Modelling of Mixed-Tailed Data with Normalizing Flows. In *Transactions on Machine Learning Research* (TMLR), 2024.
8. **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.
9. 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.

[1] Distinction allows the paper to be presented at NeurIPS, ICML, or ICLR.

Conference Articles

10. **E. Nalisnick**, S. Qian, C. Zhang, and Y. Wang. Human-AI Teaming Through the Lens of Calibration. In *Advances in Neural Information Processing Systems* (NeurIPS), 2026.
11. S. Frkovic, M. Jazbec, D. Zhang, C. A. Naesseth, I. Bogunovic, and **E. Nalisnick**. Re-evaluating Confidence Remasking in Masked Diffusion Language Models. In *Advances in Neural Information Processing Systems* (NeurIPS), 2026.
12. T. Li, J. Zhang, X. Wang, W. Jurayj, C. Jin, M. Farajtabar, **E. Nalisnick**, and Daniel Khashabi. Self-Compacting Language Model Agents. In *Advances in Neural Information Processing Systems* (NeurIPS), 2026.
13. T. X. Olausson*, M. Jazbec*, X. Wang, A. Solar-Lezama, C. A. Naesseth, S. Mandt, and **E. Nalisnick**. A Tale of Two Temperatures: Simple, Efficient, and Diverse Sampling from Diffusion Language Models. In *Proceedings of the 3rd Conference on Language Modeling* (COLM), 2026.
14. X. Wang*, A. Suresh*, A. Zhang*, R. More*, W. Jurayj, M. Farajtabar, D. Khashabi, and **E. Nalisnick**. Conformal Thinking: Risk Control for Reasoning on a Compute Budget. In *Proceedings of the 43rd International Conference on Machine Learning* (ICML), 2026.
15. A. Wynn, M. Jazbec, C. Peris, R. Khaziev, A. Liu, D. Khashabi, and **E. Nalisnick**. Controlling the Risk of Corrupted Contexts for Language Models via Early-Exiting. In *Proceedings of the 43rd International Conference on Machine Learning* (ICML), 2026.
16. H. M. Bui, M. Jazbec, **E. Nalisnick**, and A. Liu. Efficient and Uncertainty-Aware Diffusion Framework for Offline-to-Online Reinforcement Learning. In *Proceedings of the 43rd International Conference on Machine Learning* (ICML), 2026.
17. P. Kubaty, F. Szatkowski, G. Choczyński, **E. Nalisnick**, and B. Wójcik. Rethinking Calibration for Early-Exit Neural Networks. In *Proceedings of the 43rd International Conference on Machine Learning* (ICML), 2026.
18. A. Timans, T. Möllenhoff, C. A. Naesseth, M. E. Khan*, and **E. Nalisnick***. Joint Model and Data Sparsification via the Marginal Likelihood. In *Proceedings of the 43rd International Conference on Machine Learning* (ICML), 2026.
19. M. Schirmer*, M. Jazbec*, C. A. Naesseth, and **E. Nalisnick**. Monitoring Risks in Test-Time Adaptation. In *Advances in Neural Information Processing Systems* (NeurIPS), 2025.
20. A. Carbo and **E. Nalisnick**. Improving Handshape Representations for Sign Language Processing: A Graph Neural Network Approach. In *Proceedings of Empirical Methods in Natural Language Processing* (EMNLP), Main Conference Track, 2025.
21. A. Timans*, R. Verma*, **E. Nalisnick**, and C. A. Naesseth. On Continuous Monitoring of Risk Violations under Unknown Shift. In *Proceedings of the 41st Conference on Uncertainty in Artificial Intelligence* (UAI), 2025.
22. M. Jazbec, E. Wong-Toi, G. Xia, D. Zhang, **E. Nalisnick**, and S. Mandt. Generative Uncertainty in Diffusion Models. In *Proceedings of the 41st Conference on Uncertainty in Artificial Intelligence* (UAI), 2025.

23. R. Verma, V. Fischer, and **E. Nalisnick**. On Calibration in Multi-Distribution Learning. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency* (FAccT), 2025.

24. D. Tailor, A. Correia, **E. Nalisnick**, and C. Louizos. Approximating Full Conformal Prediction for Neural Network Regression with Gauss-Newton Influence. In *Proceedings of the 13th International Conference on Learning Representations* (ICLR), 2025.

25. O. Rønning, **E. Nalisnick**, C. Ley, P. Smyth, and T. Hamelryck. ELBOing Stein: Variational Bayes with Stein Mixture Inference. In *Proceedings of the 13th International Conference on Learning Representations* (ICLR), 2025.

26. A. Timans, C.-N. Straehle, K. Sakmann, C. A. Naesseth, and **E. Nalisnick**. Max-Rank: Efficient Multiple Testing for Conformal Prediction. In *Proceedings of the 28th International Conference on Artificial Intelligence and Statistics* (AISTATS), 2025.

27. U. Khurana, **E. Nalisnick**, and A. Fokkens. DefVerify: Do Hate Speech Models Reflect Their Dataset's Definition? In *Proceedings of the 31st International Conference on Computational Linguistics* (COLING), 2025.

28. M. Jazbec*, A. Timans*, T. H. Veljković, K. Sakmann, D. Zhang, C. A. Naesseth, and **E. Nalisnick**. Fast yet Safe: Early-Exiting with Risk Control. In *Advances in Neural Information Processing Systems* (NeurIPS), 2024.

29. J. U. Allingham, B. K. Mlodozeniec, S. Padhy, J. Antoran, D. Krueger, R. E. Turner, **E. Nalisnick**, J. M. Hernández-Lobato. A Generative Model of Symmetry Transformations. In *Advances in Neural Information Processing Systems* (NeurIPS), 2024.

30. U. Khurana, **E. Nalisnick**, A. Fokkens, and S. Swayamdipta. Crowd-Calibrator: Can Annotator Disagreement Inform Calibration in Subjective Tasks? In *Proceedings of the 1st Conference on Language Modeling* (COLM), 2024.

31. A. Timans, C.-N. Straehle, K. Sakmann, and **E. Nalisnick**. Adaptive Bounding Box Uncertainties via Two-Step Conformal Prediction. In *Proceedings of the 18th European Conference on Computer Vision* (ECCV), 2024. *Oral presentation (8% of accepted papers).*

32. M. Jazbec, P. Forré, S. Mandt, D. Zhang, and **E. Nalisnick**. Early-Exit Neural Networks with Nested Prediction Sets. In *Proceedings of the 40th Conference on Uncertainty in Artificial Intelligence* (UAI), 2024.

33. 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, outstanding student paper (1% of accepted papers).*

34. 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.

35. 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.

36. 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.

37. 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.

38. 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).*

39. 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.

40. 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.

41. 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.

42. **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.

43. 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.

44. **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.

45. **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.

46. **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.

47. 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.

48. **E. Nalisnick** and P. Smyth. Learning Priors for Invariance. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics* (AISTATS), 2018.

49. **E. Nalisnick** and P. Smyth. Learning Approximately Objective Priors. In *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence* (UAI), 2017.

50. **E. Nalisnick** and P. Smyth. Stick-Breaking Variational Autoencoders. In *Proceedings of the 5th International Conference on Learning Representations* (ICLR), 2017.

51. **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.

52. **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.

53. **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

54. A. Uzunoglu, M. Farajtabar, D. Khashabi, and **E. Nalisnick**. Two Faces of Missing Mass: Hallucination and Discovery in Language Models. *AI for Science*, NeurIPS 2026. *Oral presentation.*

55. R. Verma, R. Derr, C. A. Naesseth, V. Fischer, and **E. Nalisnick**. So What are Good Imprecise Forecasts? *Epistemic Intelligence in Machine Learning*, EurIPS 2025.

56. A. Suresh*, A. Zhang*, R. More*, X. Wang, W. Jurayj, B. Van Durme, **E. Nalisnick**, and D. Khashabi. Compute When Worth It: Risk Control for Reasoning on a Compute Budget. *Efficient Reasoning*, NeurIPS 2025.

57. R. Verma, V. Fischer, C. A. Naesseth, and **E. Nalisnick**. Conformal Prediction in Multi-Distribution Learning. *14th Symposium on Conformal and Probabilistic Prediction* (COPA), 2025.

58. M. Schirmer*, M. Jazbec*, C. A. Naesseth, and **E. Nalisnick**. Monitoring Risks in Test-Time Adaptation. *Test-Time Adaptation: Putting Updates to the Test!*, ICML 2025. *Oral presentation.*

59. M. Jazbec, E. Wong-Toi, G. Xia, D. Zhang, **E. Nalisnick**, and S. Mandt. Generative Uncertainty in Diffusion Models. *Quantify Uncertainty and Hallucination in Foundation Models*, ICLR 2025. *Best paper award (5% of accepted papers).*

60. X. Wang and **E. Nalisnick**. Vision-Language Models can Implicitly Quantify Aleatoric Uncertainty. *Quantify Uncertainty and Hallucination in Foundation Models*, ICLR 2025.

61. R. Verma, V. Fischer, and **E. Nalisnick**. On the Calibration of Conditional-Value-at-Risk. *Next Generation of AI Safety*, ICML 2024.

62. 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.*

63. M. Schirmer, D. Zhang, and **E. Nalisnick**. Beyond Top-Class Agreement: Using Divergences to Forecast Performance under Distribution Shift. *Distribution Shifts*, NeurIPS 2023.

64. 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.

65. S. Wang and **E. Nalisnick**. Active Learning for Multilingual Fingerspelling Corpora. *Adaptive Experimental Design and Active Learning in the Real World*, ICML 2022.

66. J. U. Allingham and **E. Nalisnick**. A Product of Experts Approach to Early-Exit Ensembles. *Dynamic Neural Networks*, ICML 2022.

67. P. Manggala, H. Hoos, and **E. Nalisnick**. Bayesian Weak Supervision via an Optimal Transport Approach. *Human-Machine Collaboration and Teaming*, ICML 2022.
68. 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.
69. 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.
70. 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.
71. P. Manggala, H. Hoos, and **E. Nalisnick**. Bayesian Regression from Multiple Sources of Weak Supervision. *Machine Learning for Data*, ICML 2021.
72. Y. Zhang and **E. Nalisnick**. On the Inconsistency of Bayesian Inference for Misspecified Neural Networks. *Symposium on Advances in Approximate Bayesian Inference*, 2021.
73. **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.
74. **E. Nalisnick** and J. M. Hernández-Lobato. Automatic Depth Determination for Bayesian ResNets. *Bayesian Deep Learning*, NeurIPS 2018.
75. 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.
76. D. Ji, **E. Nalisnick**, and P. Smyth. Mondrian Processes for Flow Cytometry Analysis. *Machine Learning for Health*, NeurIPS 2017.
77. **E. Nalisnick** and P. Smyth. Variational Inference with Stein Mixtures. *Advances in Approximate Bayesian Inference*, NIPS 2017.
78. **E. Nalisnick** and P. Smyth. The Amortized Bootstrap. *Implicit Models*, ICML 2017. *Oral presentation.*
79. **E. Nalisnick**, L. Hertel, and P. Smyth. Approximate Inference for Deep Latent Gaussian Mixtures. *Bayesian Deep Learning*, NeurIPS 2016.
80. **E. Nalisnick** and P. Smyth. Nonparametric Deep Generative Models with Stick-Breaking Priors. *Data-Efficient Machine Learning*, ICML 2016. *Oral presentation.*
81. 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. In Search of Missing Mass: A Simple Statistic for Discovery Capacity in Language Models 2026
   Trustworthy AI: From Hallucinations to Reliable Autonomy
   Simons Institute for the Theory of Computing Berkeley, USA
2. Learning to Defer to Multiple Experts: From individuals to populations and crowds 2025
   Invited Session on *The Price of Insight: Cost-aware and human-guided learning*
   Intl. Conf. on Computational and Methodological Statistics (CMStatistics) London, UK
3. Statistical Methods for AI Safety with Industrial Applications 2025
   Invited Session on *Harnessing Multi-Sector Collaboration in Statistics and Artificial Intelligence*
   Joint Statistical Meetings (JSM) Nashville, USA
4. Anomalous Anomalies: Monitoring and Adapting Anomaly Detectors 2025
   Workshop on *Anomaly Detection in Scientific Domains*
   The AAAI Conference on Artificial Intelligence Philadelphia, USA
5. Bayes Plays the Lottery: A Study of Partially Stochastic Neural Networks 2024
   Dagstuhl Seminar on Rethinking Bayes in the Age of Modern AI Wadern, Germany
6. The Boons of Being Less Bayesian: a Study of Partially Stochastic Neural Networks 2024
   Invited Session on *Advances in Inference and Theory for Bayesian Neural Networks*
   Joint Statistical Meetings (JSM) Portland, USA
7. Detecting Distribution Shift with Deep Generative Models 2023
   Dutch Society of Pattern Recognition and Image Processing Amsterdam, Netherlands
8. Towards Anytime Uncertainty Estimation in Early-Exit Neural Networks 2023
   Workshop on *Uncertainty Quantification for Computer Vision*
   International Conference on Computer Vision (ICCV) Paris, France
9. Learning to Defer to One, Multiple, or a Population of Expert(s) 2023
   ELLIS Workshop on Robust Machine Learning (RobustML) Helsinki, Finland
10. Bayesian Learning & Uncertainty Quantification 2023
    Mediterranean Machine Learning Summer School Thessaloniki, Greece
11. Towards Anytime Computation in Deep Architectures 2023
    Dagstuhl Seminar on Deep Generative Models Wadern, Germany
12. Towards Informative Priors for Bayesian Deep Learning 2022
    Dagstuhl Seminar on Tractable Probabilistic Models Wadern, Germany
13. Predictive Complexity Priors 2021
    GenU Workshop: Gen. Models and Uncertainty Quantification Copenhagen, Denmark
14. Detecting Distribution Shift with Deep Generative Models 2020
    Workshop on *Invertible Neural Nets, Normalizing Flows, and Explicit Likelihood Models* (INNF+)
    International Conference on Machine Learning (ICML) Virtual

15. Normalizing Flows for Tractable Probabilistic Modeling and Inference 2019
    T-Prime, NeurIPS Social Vancouver, Canada

Department Seminars, Companies, Local Events

1. Human-AI Teaming Through the Lens of Calibration 2026
   Univ. of Maryland: Statistics
2. Learning to Defer to One, Multiple, or a Population of Expert(s) 2023 - 2025
   Amazon, Johns Hopkins Applied Physics Laboratory, Vanderbilt Machine Learning Seminar
3. Efficient Collection of Sign Language Data 2024
   Gallaudet University
4. Towards a Statistical Foundation for Human-AI Collaboration 2023
   Univ. of Tübingen / Bosch AI, Oregon State Univ.: Computer Science, Johns Hopkins Univ.: Computer Science, George Washington Univ.: Computer Science
5. On the Calibration of Learning-to-Defer Systems 2022
   Univ. of California, Riverside: Computer Science, Univ. of Manchester: Statistics
6. Predictive Complexity Priors 2021
   Univ. of Edinburgh: ANC Seminar, Imperial College: Statistics, Alan Turing Institute
7. Detecting Distribution Shift with Deep Generative Models 2020
   Sydney Machine Learning Meetup
8. Building and Critiquing Models for Probabilistic Deep Learning 2020
   University College London: Gatsby Unit, Carnegie Mellon Univ.: Statistics, Univ. of North Carolina: Computer Science, Univ. of Amsterdam: Informatics
9. Deep Learning & Statistics: Bridging the Gap with Prob. Structure 2020
   Univ. of California, Santa Barbara: Statistics
10. Deep Learning Under Covariate Shift 2019
    Univ. of California, Irvine: AI/ML Seminar
11. Deep Learning: A Synthesis from Probabilistic Foundations 2018 - 2019
    RAND Corporation: Statistics, Univ. of California, Irvine: Statistics
12. Evaluating Deep Generative Models on Out-of-Distribution Inputs 2019
    Univ. of Oxford: Statistics, CamAIML (MSR Cambridge), Univ. of Cambridge: LTL Seminar
13. Structured Shrinkage Priors for Neural Networks 2018
    Imperial College: Statistics
14. Approximate Inference for Frequentist Uncertainty Estimation 2017
    SoCal ML Symposium
15. Deep Generative Models with Stick-Breaking Priors 2017
    Univ. of California, Irvine: AI/ML Seminar, OpenAI

General Audience, Outreach

1. A Tutor, a Tool, and a Sounding Board: How Statistics Students Use Generative AI 2026
   Chance Magazine, co-authored with M. Glickman, J. Yan, and M. Çetinkaya-Rundel
2. AI for Sign Language Technologies, Maryland School for the Deaf 2025
3. Discussant for the ASA webinar *Statistical Foundations to Survive the AI Apocalypse* 2025
4. Discussed Human-AI Collaboration on the *Faculty Factory Podcast* 2024
5. AI Risks: From Today to Doomsday — An Academic Panel Discussion 2023
   Amsterdam ELLIS Unit