Stephan Rabanser

Last update: December 3, 2025

✉ rabanser@princeton.edu
https://rabanser.dev

Experience

- **Postdoctoral Research Fellow** — Princeton, NJ
  Princeton University, advised by Prof. Arvind Narayanan & Prof. Matthew Salganik — *Since October 2025*
- **Machine Learning Researcher** — Toronto, Canada
  Vector Institute for Artificial Intelligence — *September 2020 – August 2025*
- **Student Researcher** — Zurich, Switzerland
  Google Research — *August 2024 – January 2025*
  - Developed hierarchical selective prediction/rejection techniques for large vision-language models (VLMs).
- **Intern Applied Scientist** — Munich, Germany
  Amazon, AWS AI Labs — *June 2021 – October 2021*
  - Designed context-invariant time series representations using contrastive and domain-adversarial learning.
- **Intern Applied Scientist** — Munich, Germany
  Amazon, AWS AI Labs — *September 2019 – July 2020*
  - Systematically assessed the impact of I/O representations for deep-learning-based time-series forecasting.

Education

- **PhD in Computer Science** — Toronto, Canada
  University of Toronto, advised by Prof. Nicolas Papernot — *September 2020 – August 2025*
  - **Supervisory Committee**: Prof. Nicolas Papernot, Prof. Rahul Krishnan, Prof. David Duvenaud, Prof. Roger Grosse, Prof. Zachary Lipton
  - **Research Interests**: Machine Learning, Robustness, Safety, Reliability, Uncertainty, Causality, Generative Modeling, Representation Learning, Probabilistic Deep Learning, Anomaly Detection, Distribution Shifts, Out-of-Distribution Sample Detection.
- **Visiting Graduate Student** — Cambridge, UK
  University of Cambridge, advised by Prof. David Krueger — *June 2023 – September 2023*
- **M.Sc. in Computer Science** — Munich, Germany
  Technical University of Munich (TUM), advised by Prof. Stephan Günnemann — *October 2015 – July 2019*
- **Visiting Research Scholar** — Pittsburgh, PA
  Carnegie Mellon University (CMU), advised by Prof. Zachary Lipton — *August 2018 – January 2019*
- **Honours Degree in Technology Management** — Munich, Germany
  Center for Digital Technology and Management (CDTM) — *August 2015 – June 2017*
- **Visiting Research Student** — Cambridge, MA
  Massachusetts Institute of Technology (MIT), advised by Prof. Thomas Malone — *February 2016 – June 2016*
- **B.Sc. in Computer Science, Minor in Economic Sciences** — Munich, Germany
  Technical University of Munich (TUM) — *October 2012 – October 2015*

Awards & Honors

- **Top Reviewer Award** — *ICML 2025, NeurIPS 2023, Dist. Shift Workshop @ NeurIPS 2021*
- **Member of the Elite Network of Bavaria** — *Since 2016*
- **Apple WWDC Student Scholarship** — *June 2013*

Publications

Stephan Rabanser and Nicolas Papernot. **What Does It Take to Build a Performant Selective Classifier?** In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. [Paper]

Youhe Jiang, Fangcheng Fu, Wanru Zhao, Stephan Rabanser, Nicholas D Lane, and Binhang Yuan. **Cascadia: A Cascade Serving System for Large Language Models**. *arXiv preprint arXiv:2506.04203*, 2025. [Paper]

Stephan Rabanser, Nathalie Rauschmayr, Achin Kulshrestha, Petra Poklukar, Wittawat Jitkrittum, Sean Augenstein, Congchao Wang, and Federico Tombari. **Gatekeeper: Improving Model Cascades Through Confidence Tuning**. *Advances in Neural Information Processing Systems (NeurIPS) & TTODLer-FM Workshop @ ICML (**Best Poster Award**)*, 2025. [Paper]

Angline Pouget, Mohammad Yaghini, Stephan Rabanser, and Nicolas Papernot. **Suitability Filter: A Statistical Framework for Model Evaluation in Real-World Deployment Settings**. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2025. **Oral presentation**. [Paper]

Stephan Rabanser, Ali Shamsabadi, Olive Franzese, Xiao Wang, Adrian Weller, and Nicolas Papernot. **Confidential Guardian: Cryptographically Prohibiting the Abuse of Model Abstention**. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2025. [Paper]

Stephan Rabanser, Anvith Thudi, Kimia Hamidieh, Adam Dziedzic, and Nicolas Papernot. **Selective Prediction Via Training Dynamics**. *Transactions on Machine Learning Research*, 2025. [Paper]

Stephan Rabanser, Anvith Thudi, Abhradeep Thakurta, Krishnamurthy Dvijotham, and Nicolas Papernot. **Training Private Models That Know What They Don't Know**. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. [Paper, Slides]

Nicholas Franzese, Adam Dziedzic, Christopher A. Choquette-Choo, Mark R. Thomas, Muhammad Ahmad Kaleem, Stephan Rabanser, Congyu Fang, Somesh Jha, Nicolas Papernot, and Xiao Wang. **Robust and Actively Secure Collaborative Machine Learning**. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. [Paper]

Adam Dziedzic, Stephan Rabanser, Mohammad Yaghini, and Nicolas Papernot. ***p*-DkNN: Out-of-Distribution Detection through Statistical Testing of Deep Representations**. *arXiv preprint arXiv:2207.12545*, 2022. [Paper]

Stephan Rabanser, Tim Januschowski, Kashif Rasul, Oliver Borchert, Richard Kurle, Jan Gasthaus, Michael Bohlke-Schneider, Nicolas Papernot, and Valentin Flunkert. **Intrinsic Anomaly Detection in Multi-Variate Time Series**. *arXiv preprint arXiv:2206.14342*, 2022. [Paper]

Stephan Rabanser, Tim Januschowski, Valentin Flunkert, David Salinas, and Jan Gasthaus. **The Effectiveness of Discretization in Forecasting: An Empirical Study on Neural Time Series Models**. In *7th KDD Workshop on Mining and Learning from Time Series (MiLeTS)*, 2020. **Oral presentation**. [Paper, Slides]

Stephan Rabanser, Stephan Günnemann, and Zachary Lipton. **Failing Loudly: An Empirical Study of Methods for Detecting Dataset Shift**. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. [Paper, Poster, Slides]

Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. **Introduction to Tensor Decompositions and their Applications in Machine Learning**. *arXiv preprint arXiv:1711.10781*, 2017. [Paper]

Community Service

- **Reviewing**: NeurIPS (2025, 2024, 2023, 2022, 2021), ICML (2025, 2022, 2021), ICLR (2024), IEEE SatML (2024), Distribution Shift Workshop @ ICML (2022), Distribution Shift Workshop @ NeurIPS (2023, 2022, 2021), Human Evaluation of Generative Models Workshop @ NeurIPS (2022), Time Series Workshop @ ICML (2021), Time Series Workshop @ KDD (2022), AAAI (2020)
- **Talks**: Google DeepMind London (Sep 2023), MIT MIMO Student Research Forum (Oct 2022), Intel Private AI Institute Fall Workshop (Oct 2022), Microsoft Security Data Science Colloquium (Jul 2021)