

Rolf Jagerman

rjagerman@gmail.com ✉
www.jagerman.nl 🌐
[rolfjagerman](#) 🐦
[rjagerman](#) 🌐

Education

- 2016–2020 **PhD Computer Science**, *University of Amsterdam*.
Thesis: Efficient, Safe and Adaptive Learning from User Interactions.
- 2014–2016 **MSc Computer Science**, *ETH Zürich*, 5.52/6.0.
Thesis: Web-scale Topic Models in Spark: an Asynchronous Parameter Server.
- 2011–2014 **BSc Computer Science**, *TU Delft*, 8.7/10.0 (Cum Laude).
Thesis: Android Tor Tribler Tunneling (AT3).
Successfully completed the honours programme.

Professional Experience

- 2024–present **Senior Software Engineer**, *Google DeepMind*.
Working at the GenAI unit on Large Language Models (LLMs) + Ranking.
- 2022–2024 **Senior Software Engineer**, *Google*.
Research and development of Large Language Models for Ranking. Driven impact across Google Cloud AI, YouTube, Chrome Web Store, and more.
- 2020–2022 **Software Engineer**, *Google*.
Research and development of Learning-to-Rank techniques. Impactful launches across YouTube, Chrome Web Store, and more.
- 2020 **Machine Learning Intern**, *Apple*.
Winter internship at Apple.
- 2019 **Software Engineering Intern**, *Google*.
Summer research internship at Google.
- 2015 **Research assistant**, *Data Analytics Lab*, *ETH Zürich*.
Developed a large-scale distributed implementation of the Latent Dirichlet Allocation (LDA) model in Spark to run on a 20+ TB dataset.
- 2008–2020 **Founder**, *Contended*.
Freelance web design and software development for various companies in the Netherlands.

Publications

- 2024 H. Oosterhuis, **R. Jagerman**, Z. Qin, X. Wang and M. Bendersky. Reliable Confidence Intervals for Information Retrieval Evaluation Using Generative A.I.. *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, 2024.

- Z. Qin, **R. Jagerman**, K. Hui, H. Zhuang, J. Wu, L. Yan, J. Shen, T. Liu, J. Liu, D. Metzler, X. Wang and M. Bendersky. Large Language Models are Effective Text Rankers with Pairwise Ranking Prompting. *Proceedings of the 2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. ACL, 2024. [PDF]
- 2023 Z. Qin, **R. Jagerman**, R. Pasumarthi, H. Zhuang, H. Zhang, A. Bai, K. Hui, L. Yan and X. Wang. RD-Suite: A Benchmark for Ranking Distillation. *Proceedings of the 37th Conference on Neural Information Processing Systems*. NeurIPS, 2023. [PDF]
- R. Jagerman**, H. Zhuang, Z. Qin, X. Wang and M. Bendersky. Query Expansion by Prompting Large Language Models. *Gen-IR Workshop at SIGIR 2023*. arXiv, 2023. [PDF]
- 2022 A. Bai, **R. Jagerman**, Z. Qin, P. Kar, B. Lin, X. Wang, M. Bendersky and M. Najork. Regression Compatible Listwise Objectives for Calibrated Ranking. *arXiv*. arXiv, 2022. [PDF]
- H. Zhuang, Z. Qin, **R. Jagerman**, K. Hui, J. Ma, J. Lu, J. Ni, X. Wang and M. Bendersky. RankT5: Fine-Tuning T5 for Text Ranking with Ranking Losses. *arXiv*. arXiv, 2022. [PDF]
- R. Jagerman**, X. Wang, H. Zhuang, Z. Qin, M. Bendersky and M. Najork. Rax: Composable Learning-to-Rank using JAX. *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, 2022. [PDF]
- R. Jagerman**, Z. Qin, X. Wang, M. Bendersky and M. Najork. On Optimizing Top-K Metrics for Neural Ranking Models. *Proceedings of the 45th ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2022. [PDF]
- 2021 Z. Qin, H. Zhuang, **R. Jagerman**, X. Qian, P. Hu, C. Chen, X. Wang, M. Bendersky and M. Najork. Bootstrapping Recommendations at Chrome Web Store. *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2021. [PDF]
- R. Jagerman**, W. Kong, R. Pasumarthi, Z. Qin, M. Bendersky and M. Najork. Improving Cloud Storage Search with User Activity. *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. ACM, 2021. [PDF]
- 2020 **R. Jagerman**. Efficient, Safe and Adaptive Learning from User Interactions. *PhD Thesis*, 2020. [PDF]
- R. Jagerman** and M. de Rijke. Accelerated Convergence for Counterfactual Learning to Rank. *Proceedings of the 43rd ACM International Conference on Research and Development in Information Retrieval*. ACM, 2020. [PDF]
- R. Jagerman**, I. Markov and M. de Rijke. Safe Exploration for Optimizing Contextual Bandits. *Transactions on Information Systems*. ACM, 2020. [PDF]

- 2019 **R. Jagerman**, H. Oosterhuis and M. de Rijke. To Model or to Intervene: A Comparison of Counterfactual and Online Learning to Rank from User Interactions. *Proceedings of the 42nd ACM International Conference on Research and Development in Information Retrieval*. ACM, 2019. [PDF]
- R. Jagerman**, I. Markov and M. de Rijke. When People Change their Mind: Off-Policy Evaluation in Non-stationary Recommendation Environments. *Proceedings of the 12th ACM International Conference on Web Search and Data Mining*. ACM, 2019. [PDF]
- 2018 **R. Jagerman**, K. Balog and M. de Rijke. OpenSearch: Lessons Learned from an Online Evaluation Campaign. *Journal of Data and Information Quality*. ACM, 2018. [PDF]
- 2017 **R. Jagerman**, K. Balog, P. Schaer, J. Schaible, N. Tavakolpoursaleh and M. de Rijke. Overview of TREC OpenSearch 2017. *Proceedings of the 26th Text REtrieval Conference*, 2017. [PDF]
- R. Jagerman**, H. Oosterhuis and M. de Rijke. Query-level Ranker Specialization. *1st International Workshop on LEARning Next gERneration Rankers (LEARNER)*, 2017. [PDF]
- R. Jagerman**, J. Kiseleva and M. de Rijke. Modeling Label Ambiguity for Neural List-Wise Learning to Rank. *2nd International Workshop on Neural Information Retrieval (Neu-IR)*, 2017. [PDF]
- R. Jagerman**, C. Eickhoff and M. de Rijke. Computing Web-scale Topic Models using an Asynchronous Parameter Server. *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2017. [PDF]

Skills

Languages	Dutch (native speaker), English (fluent)
Programming	Python, Scala, Java, C++
Frameworks & Tools	Numpy, Scipy, JAX, Numba, Scikit-Learn, Tensorflow, PyTorch, Spark

Open Source Contributions

Rax	Composable Learning-to-Rank using JAX.
PyTorchLTR	Learning-to-Rank in PyTorch.
Glint	Spark-compatible parameter server.
Shoelace	Neural List-wise Learning-to-Rank library for Chainer.
Murmur3.jl	High performance Julia implementation of Murmur3 hashing.
Django	Bug fixes in the administration panel UI JavaScript.
Tribler	Bug fixes for software stability on Android.