

Curriculum Vitae - Korbinian Pöppel

Name Korbinian Pöppel (PhD)
Email korbi@korbi.ai
Github <https://github.com/kpoeppel>
Twitter [@KorbiPoeppel](https://twitter.com/KorbiPoeppel)
Website <https://korbi.ai>



Education

2012 – 2015 Bachelor Physics, LMU Munich
grade: 1.02, 'best of semester'
2015 – 2018 Elite-Master Theoretical and Mathematical Physics, LMU/TUM,
grade: 1.1, 'with great distinction'
Master Physics, LMU Munich, grade: 1.1
2019 – 2023 Advanced Studies in Informatics (TU Munich) – grade: 1.2
2023 – 2025 ELLIS PhD Student JKU Linz – Artificial Intelligence – grade: 1.0
(best possible)

Scholarship

2025 Konrad Zuse Visiting ELLIS PhD Student Scholarship

Positions

2016 Research Assistant, LS Frey LMU Munich – Evolutionary Systems
2018 Research Assistant, LS Frey LMU Munich – Pattern Formation
2018 – Self-Employment in Software Development
2018 – 2019 PhD Student at MPI-IS Stuttgart (Active Acoustic Holography)
Advisor: Prof. Peer Fischer – resignation on own decision
Associate of Center of Learning Systems
2020 – 2021 Thesis: „Explainable NLP with Knowledge Graphs“ at inovex GmbH
2020 – 2021 Co-Founder of Norona GbR (Making Rapid tests visible)
2021 – 2022 First Machine Learning Engineer at docsy GmbH (now Knowlix)

- built OCR evaluation pipeline, multi-modal pre-training
- lead data augmentation pipeline, enabling privacy preservation
- multi-modal training for document classification and NER

2023 Research Engineer at IARAI GmbH (in transition to JKU)
2023 – 2025 Research Assistant at JKU Linz, publ.: ICML, ICLR, NeurIPS
Advisor: Prof. Sepp Hochreiter

- co-developing xLSTM architecture (xLSTM), numerical optimization, testing capabilities on different synthetic data
- first to introduce kernel development (FlashRNN) in CUDA and Triton at the lab, first fused open RNN kernels with automatic shape adaptation for kernel fusion
- extending RNNs stably to multiple dimensions and graphs (pLSTM), new benchmark for multi-dimensional length extrapolation

2024 – 2025 Co-Founder of OpenPV GbR (openpv.de)

- prototyping and feature development for web application

2024 – 2025 PhD Researcher NXAI GmbH

- scaling to xLSTM-7B LLM on 128 GPUs, fastest 7B model
- leading LLM evaluation and long context extension training

2025 – Post-Doctoral Research Engineer ELLIS Institute Tübingen gGmbH

- unifying experimentation infrastructure for a team of 50+
- scaling and speed optimization on big European clusters up to 2048 GPUs for dense and MoE models

Publications

Shulgin, E., Franke, J. K. H., von Rütte, D. et al. (2026). Deriving Hyperparameter Scaling Laws via Modern Optimization Theory. *High-Dimensional Learning Dynamics 2026*. <https://openreview.net/forum?id=gN4illcXOd>.

Siems, J., Grazzi, R., **Pöppel, K.**, Kalinin, K., Ballani, H., and Rahmani, B. (2026). Learning State-Tracking from Code Using Linear RNNs. <https://arxiv.org/abs/2602.14814>.

Shakya, S., Hartl, A. R., Hochreiter, S., & **Pöppel, K.** (2025) Adaptive Retrieval helps Reasoning in LLMs-but mostly if it's not used. In *EurIPS 2025 Workshop on Principles of Generative Modeling (PriGM)*. Retrieved from <https://openreview.net/pdf?id=lldcyeuR2r>

Pöppel, K. and Freinschlag, R. and Schmied, T. and Lin, W. and Hochreiter, S. (2025). pLSTM: parallelizable Linear Source Transition Mark networks. The Thirty-ninth Annual Conference on Neural Information Processing Systems. Retrieved from <https://openreview.net/forum?id=2sa13vyCn0>

Beck, M. and **Pöppel, K.** and Lippe, P., Hochreiter, S. (2025). Tiled Flash Linear Attention: More Efficient Linear RNN and xLSTM Kernels. The Thirty-ninth Annual Conference on Neural Information Processing Systems. Retrieved from <https://openreview.net/forum?id=b6H64u6Tql>

Beck, M.* , **Pöppel, K.***, Lippe, P.* Kurle, R., Blies, P. M., Klambauer, G., Böck, S. and Hochreiter, S. (2025). xLSTM 7B: A Recurrent LLM for Fast and Efficient Inference. Forty-second conference on Machine Learning. Retrieved from <https://openreview.net/forum?id=LV3DpKD08B>

Schmied, T., Adler, T., Patil, V., Beck, M., **Pöppel, K.**, Brandstetter, J., Klambauer, G., Pascanu, R., & Hochreiter, S. (2025). A large recurrent action model: xLSTM enables fast inference for robotics tasks. Forty-second International Conference on Machine Learning. Retrieved from <https://openreview.net/forum?id=J2JxJ0P1LI>

Pöppel, K., Beck, M., Hochreiter S. (2025). FlashRNN: I/O-Aware Optimization of Traditional RNNs on Modern Hardware. The Thirteenth International Conference on Learning Representations. Retrieved from <https://openreview.net/forum?id=l0ZzTvPfTw>

Alkin, B., Beck, M., **Pöppel, K.**, Hochreiter, S., & Brandstetter, J. (2024). Vision-LSTM: xLSTM as generic vision backbone. The Thirteenth International Conference on Learning Representations. Retrieved from <https://openreview.net/forum?id=SiH7DwNKZZ>

Beck, M.* , **Pöppel, K.***, Spanring, M., Auer, A., Prudnikova, O., Kopp, M. K., Klambauer, G., Brandstetter, J., & Hochreiter, S. (2024). xLSTM: Extended long short-term memory. *The Thirty-Eighth Annual Conference on Neural Information Processing Systems*. Retrieved from <https://openreview.net/forum?id=ARAxPPIAhq>

Ma, Z., Holle, A. W., Melde, K., Qiu, T., **Pöppel, K.**, Kadiri, V. M., & Fischer, P. (2020). Acoustic holographic cell patterning in a biocompatible hydrogel. *Advanced Materials*, 32(4), 1904181. <https://doi.org/10.1002/adma.201904181>