Maurice Weber

AI Researcher

mauriceweber.github.io | Maurice Weber | mauriceweber | maurice@together.ai

Education

ETH Zurich Ph.D. in Computer Science
February 2020 – Current | Zurich, Switzerland
Exploring robustness guarantees for machine learning systems and quantum algorithms, in the context of adversarial attacks, naturally occurring noise and distribution shifts.
Supervisors: Prof. Ce Zhang & Prof. Martin Vechev

ETH Zurich MSc Mathematics (graduated with distinction)
September 2016 – May 2019 | Zurich, Switzerland
Coursework with focus on machine learning and mathematical statistics. Master thesis on lossy image compression with recurrent neural networks.
Supervisors: Prof. Ce Zhang & Prof. Nicolai Meinshausen.
GPA: 5.77 / 6.0

ETH Zurich BSc Mathematics
September 2011 – June 2015 | Zurich, Switzerland
Courses covered wide areas of Mathematics including Linear Algebra, Functional Analysis, Topology, Probability Theory, Brownian Motion & Stochastic Calculus. Bachelor thesis on value at risk with defaultable securities.
GPA: 4.85 / 6.0

Experience

Together AI | AI Researcher
January 2024 - present | Zurich, Switzerland
- Member of the Research team at Together AI, working on research related to the training data of large language models.

Xanadu Quantum Technologies | Research Intern
May 2022 - August 2022 | Toronto, Canada
- Made central contributions to a research project on using transformer-based generative models for quantum state tomography. This led to a research paper which we submit to one of the top Physics journals.
- Contributed code and code reviews to the PennyLane software library.

ETH Zurich | Teaching & Research Assistant
April 2019 - Current | Zurich, Switzerland
- Teaching exercise classes and supporting exams for courses offered by the department of computer science.

Pictet Asset Management | Intern Multi-Asset Investments (part-time)
May 2016 - May 2018 | Zurich, Switzerland
- Assisted the multi asset teams in Geneva and Zurich in their daily investment process and developed tools for analysis of portfolio exposure to currencies, regions and industries.

UBS Investment Bank | Intern Equity Derivatives Trading
May 2015 - April 2016 | Zurich, Switzerland
- Assisted the derivatives trading teams in their daily routines.
- Developed trade supporting tools using MS Visual Basic, including a tool for real-time monitoring of competitor prices.

Profile

I am an AI researcher focusing on aspects of large language models related to pre-training data. Prior to my role at Together AI, I obtained a Ph.D. in student in Computer Science from ETH Zurich where I have published scientific research in the area of reliable and trustworthy ML at top conferences and journals.

Skills

Python | Java | SQL | git | LaTeX | Bash | R | Big Data | Communication | Teamwork | Organization | Creativity

Expertise

- Scientific Research
- Software Development
- Machine Learning
- Large Language Models

Languages

German: ●●●●● Native
English: ●●●●○ Fluent
Spanish: ●●●●○ Fluent
French: ●●●○○ Intermediate

Open Source

- Core contributions to the RedPajama datasets (collaboration with Together.ai).
- Contributed transpiler module and code reviews to PennyLane quantum software library.
- Contributed robustness certification application to the Tequila quantum software library.

Service

Reviewer Physical Review A, Physical Review Research, PRX Quantum, npj Quantum Information, IEEE Transactions on Neural Networks and Learning Systems.

Publications & Preprints

(* denotes equal contribution)

2023 **Maurice Weber***, Carlo Siebenschuh*, Rory Marshall Butler*, Anton Alexandrov, Valdemar Ragnar Thanner, Georgios Tsolakis, Haris Jabbar, Ian Foster, Bo Li, Rick Stevens, and Ce Zhang. Wordscape: a pipeline to extract multilingual, visually rich documents with layout annotations from web crawl data. In Advances in Neural Information Processing Systems, 2023

2023 **Maurice Weber***, Xiaojun Xu*, Bojan Karlaš, Ce Zhang, and Bo Li. Rab: Provable robustness against backdoor attacks. In 2023 IEEE Symposium on Security and Privacy (SP), 2023

2022 Haoxiang Wang*, **Maurice Weber***, Josh Izaac, and Cedric Yen-Yu Lin. Predicting properties of quantum systems with conditional generative models. arXiv preprint arXiv:2211.16943, 2022

2022 **Maurice Weber**, Linyi Li, Boxin Wang, Zhikuan Zhao, Bo Li, and Ce Zhang. Certifying out-of-domain generalization for blackbox functions. In 39th International Conference on Machine Learning, 2022

2022 **Maurice Weber**, Abhinav Anand, Alba Cervera-Lierta, Jakob S Kottmann, Thi Ha Kyaw, Bo Li, Alán Aspuru-Guzik, Ce Zhang, and Zhikuan Zhao. Toward reliability in the nisq era: Robust interval guarantee for quantum measurements on approximate states. Physical Review Research, 4(3):033217, 2022

2022 Nicolas Langer, **Maurice Weber**, Bruno Hebling Vieira, Dawid Strzelczyk, Lukas Wolf, Andreas Pedroni, Jonathan Heitz, Christoph Schultheis, et al. The ai neuropsychologist: Automatic scoring of memory deficits with deep learning. bioRxiv, 2022

2022 Mintong Kang*, Linyi Li*, **Maurice Weber**, Yang Liu, Ce Zhang, and Bo Li. Certifying some distributional fairness with subpopulation decomposition. In Advances in Neural Information Processing Systems 35 (NeurIPS 2022), 2022

2021 **Maurice Weber**, Nana Liu, Bo Li, Ce Zhang, and Zhikuan Zhao. Optimal provable robustness of quantum classification via quantum hypothesis testing. npj Quantum Information, 7(1):1–12, 2021

2021 Linyi Li*, **Maurice Weber***, Xiaojun Xu, Luka Rimanic, Bhavya Kailkhura, Tao Xie, Ce Zhang, and Bo Li. Tss: Transformation-specific smoothing for robustness certification. In 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21), 2021

2020 **Maurice Weber**, Cedric Renggli, Helmut Grabner, and Ce Zhang. Observer dependent lossy image compression. In 42nd German Conference on Pattern Recognition, 2020

2019 Filipe Barata, Kevin Kipfer, **Maurice Weber**, Peter Tinschert, Elgar Fleisch, and Tobias Kowatsch. Towards device-agnostic mobile cough detection with convolutional neural networks. In 2019 IEEE International Conference on Healthcare Informatics (ICHI), pages 1–11. IEEE, 2019

Conferences & Workshops

May 2023 **IEEE Symposium on Security and Privacy (SP) 2023** | San Francisco CA, US
Talk: RAB: Provable Robustness against Backdoor Attacks

July 2022 **39th International Conference on Machine Learning (ICML) 2022** | Baltimore MD, US
Spotlight Talk: Certifying Out-of-Domain Generalization for Blackbox Functions

Mar 2022 **26th Conference on Quantum Information Processing (QIP) 2022** | Los Angeles CA, US
Poster: Toward Reliability in the NISQ Era: Robust Interval Guarantee for Quantum Measurements on Approximate States

Feb 2021 **25th Conference on Quantum Information Processing (QIP) 2021** | Virtual
Poster: Optimal Provable Robustness of Quantum Classification via Quantum Hypothesis Testing

Dec 2020 **20th Asian Quantum Information Science Conference (AQIS) 2021** | Virtual
Talk: Optimal Provable Robustness of Quantum Classification via Quantum Hypothesis Testing

Sep 2020 **42nd German Conference on Pattern Recognition (GCPR) 2020** | Virtual
Contributed Talk: Observer Dependent Lossy Image Compression

July 2020 **ICML 2020 Workshop on Uncertainty & Robustness in Deep Learning** | Virtual
Poster: Provable Robust Learning Based on Transformation-Specific Smoothing

Teaching

Fall 2022 **Computer Science (D-MATH/D-PHYS ETH Zurich)** | Teaching Assistant

Spring 2022 **Data Modelling and Databases (D-INFK ETH Zurich)** | Teaching Assistant

Fall 2020 **Information Systems for Engineers (D-INFK ETH Zurich)** | Teaching Assistant

Spring 2017 **Linear Algebra II (D-MATH ETH Zurich)** | Teaching Assistant

Spring 2015 **Calculus II (D-MATH ETH Zurich)** | Teaching Assistant

References

PROF. CE ZHANG

ce@together.ai
CTO at Together AI; Associate Professor at University of Chicago.

Additional references available upon request.