

TIANLIN LIU

✉ t.liu@unibas.ch 🏠 tianlinliu.com 🌐 liutianlin0121

Last updated: June 8, 2024

EXPERIENCE

- Fall 2023 **Research Scientist Intern at Google DeepMind, Paris.**
- Developed novel reinforcement learning from human feedback (RLHF) techniques for aligning language models [13].
 - Hosted by Mathieu Blondel, Felipe Llinares, and Daniele Calandriello.
- Summer 2022 **Research Scientist Intern at Google Brain, Zurich.**
- Investigated optimal transport approaches for mixture-of-experts (MoE) models for computer vision tasks [9, 12].
 - Hosted by Joan Puigcerver and Mathieu Blondel.

EDUCATION

- 2020–PRESENT **Ph.D. Computer Science**
University of Basel, Switzerland
Advisor: Ivan Dokmanić
- 2017–2019 **M.S. Data Engineering**
Jacobs University Bremen, Germany
Advisor: Herbert Jaeger
- 2013–2016 **B.S. Mathematics**
Jacobs University Bremen, Germany
Advisor: Götz Pfander

AWARDS AND HONORS

- 2023 Google Conference Scholarship
- 2022 Young Researcher, Heidelberg Laureate Forum.
- 2022 Travel Award, SIAM Conference on Imaging Science.
- 2021 Oberwolfach Leibniz Graduate Student Grant, Oberwolfach Institute.
- 2019 Dean's Award for Outstanding Master's Thesis [4], Jacobs Univ. Bremen.
- 2019 IEEE NER 2019 Best Paper Finalist Award [2].
- 2018 SmartStart-1 Fellowship, Bernstein Network and Volkswagen Foundation.

PUBLICATIONS

Google Scholar

(★ indicates equal contributions)

- [13] **Decoding-time Realignment of Language Models.**
Tianlin Liu, Shangmin Guo, Leonardo Bianco, Daniele Calandriello, Quentin Berthet, Felipe Llinares, Jessica Hoffmann, Lucas Dixon, Michal Valko, Mathieu Blondel.
In Proc. of the International Conference on Machine Learning (ICML), 2024.
Spotlight presentation
- [12] **Routers in Vision Mixture of Experts: An Empirical Study.**
Tianlin Liu, Mathieu Blondel, Carlos Riquelme, Joan Puigcerver.
Transactions on Machine Learning Research (TMLR), 2024.
- [11] **The N Implementation Details of RLHF with PPO.**
Shengyi Huang, **Tianlin Liu**, Leandro Von Werra.
The Blogpost Track of the International Conference on Learning Representations (ICLR), 2024.
Spotlight presentation
- [10] **WaveBench: Benchmarking Data-driven Solvers for Linear Wave Propagation PDEs.**
Tianlin Liu★, Jose Antonio Lara Benitez★, Florian Faucher★, AmirEhsan Khorashadizadeh, Maarten V. de Hoop, Ivan Dokmanić.
Transactions on Machine Learning Research (TMLR), 2024.
- [9] **Sparsity-Constrained Optimal Transport.**
Tianlin Liu, Joan Puigcerver, and Mathieu Blondel.
In Proc. of the International Conference on Learning Representations (ICLR), 2023.
Spotlight presentation
- [8] **Learning Multiscale Convolutional Dictionaries for Image Reconstruction.**
Tianlin Liu, Anadi Chaman, David Belius, and Ivan Dokmanić.
IEEE Transactions on Computational Imaging, 2022.
- [7] **Universal Approximation Under Constraints is Possible with Transformers.**
Anastasis Kratsios, Behnoosh Zamanlooy, **Tianlin Liu**, and Ivan Dokmanić.
In Proc. of the International Conference on Learning Representations (ICLR), 2022.
Spotlight presentation
- [6] **Finding Trainable Sparse Networks through Neural Tangent Transfer.**
Tianlin Liu and Friedemann Zenke.
In Proc. of the International Conference on Machine Learning (ICML), 2020.
- [5] **Causally Denoise Word Embeddings Using Half-Sibling Regression.**
Zekun Yang★ and **Tianlin Liu**★.
In Proc. of the AAAI Conference on Artificial Intelligence (AAAI), 2020.
- [4] **Harnessing Slow Dynamics in Neuromorphic Computation.**
Tianlin Liu. Master’s thesis, 2019.
Dean’s Prize for Outstanding Master’s Thesis
- [3] **Continual Learning for Sentence Representations Using Conceptors.**
Tianlin Liu, Lyle Ungar, and João Sedoc.
In Proc. of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2019.
- [2] **Reservoir Transfer on Analog Neuromorphic Hardware.**
Xu He, **Tianlin Liu**, Fatemeh Hadaeghi, and Herbert Jaeger.
In Proc. of the International IEEE Conference on Neural Engineering (NER), 2019.
Best Paper Finalist Award (3rd place out of 467 accepted papers)

- [1] **Unsupervised Post-processing of Word Vectors via Conceptor Negation.**
Tianlin Liu, Lyle Ungar, and João Sedoc.
 In Proc. of the AAAI Conference on Artificial Intelligence (AAAI), 2019.

OPEN-SOURCE CONTRIBUTION

huggingface/transformers : transformers library	☆ 120k
huggingface/trl : lightweight library for RLHF	☆ 7.3k
PythonOT/POT : python optimal transport	☆ 2.2k
google-research/vmoe : vision mixture of experts models	☆ 465
google/jaxopt : differentiable optimizers in JAX	☆ 864

JOURNAL AND CONFERENCE REVIEWING

CONFERENCE	ICML (2021–2024), NeurIPS (2021–2023), ICLR (2022, 2023), LoG (2022)
JOURNAL	IEEE Transactions on Signal Processing, IEEE Signal Processing Letters, IEEE Open Journal of Signal Processing

TEACHING EXPERIENCE

2022	Mathematics of Data Science	Uni Basel
	Scientific Writing	Uni Basel
	Machine Learning on Graphs, Groups, and Manifolds	Uni Basel
2021	Computational Sustainability and AI for Social Good	Uni Basel
2021	High-dimensional Data Analysis and Learning on Graphs	Uni Basel
2018, 2019	Machine Learning	Jacobs Uni Bremen
2018, 2019	Principles of Statistical Modeling	Jacobs Uni Bremen