

Phillip Lippe

Research Scientist working on large-scale multimodal LLMs and multimodal generation.

Nationality German
🔗 phlippe.github.io

Languages German (native), English (fluent)
@ phillip.lippe@gmail.com

📍 Mountain View, US
in [in/phillip-lippe](https://www.linkedin.com/in/phillip-lippe)

🔗 [phlippe](#) 📄 [List of Publications](#)

EXPERIENCE

Senior Research Scientist

Google DeepMind

📅 Jan 2025 – present 📍 Mountain View, US

- Working on multimodality and RL for [Gemini](#).
- Core contributor to [Gemini 2.5 Flash Image \(Nano Banana\)](#), [Imagen 4](#), [Gemini 2.5](#), [Gemini 2.0 Flash Image](#)

Senior Research Scientist

NXAI

📅 Sep 2024 – Dec 2024 📍 Amsterdam, Netherlands

- Research on scaling [xLSTM](#) as Transformer alternative.
- Training 7B parameter models over 256 H100 GPUs.
- Developing Pallas and Triton Kernels on H100 GPUs.

Student Researcher

Google DeepMind

📅 Aug 2023 – Nov 2023 📍 Amsterdam, Netherlands

- Research in generative multimodal pretraining for image-text sequences; hosted by [Mostafa Dehghani](#).
- Training up to 4B parameter models over 512 devices.
- Results internally adopted in the [Gemini](#) project.

Research Intern

Microsoft Research

📅 Mar 2023 – May 2023 📍 Amsterdam, Netherlands

- Research in neural PDE solvers for scientific simulation.
- Published *PDE-Refiner: Achieving Accurate Long Rollouts with Neural PDE Solvers* at NeurIPS 2023 ([pdf](#)).

PROJECTS

Google Developer Expert, Machine Learning

- Been an ML GDE for JAX in Aug 2022 - Dec 2024.
- Presented [tutorials](#) on distributed training for scaling deep learning models to billion parameters.

Tutorials

- Created and taught various Deep Learning tutorial notebooks in PyTorch and JAX ([website](#)).

Lecturer and Teaching Assistant

- Given lectures in the graduate courses Deep Learning, Foundation Models, and Advanced NLP at the University of Amsterdam (2019-2024).

EDUCATION

PhD, Artificial Intelligence

University of Amsterdam, QUVA lab

📅 Sep 2020 – Aug 2024 📍 Amsterdam, Netherlands

- PhD Advisors: [Efstratios Gavves](#), [Taco Cohen](#), [Sara Magliacane](#), [Yuki Asano](#).
- ELLIS PhD; collaboration with Qualcomm AI Research.
- Graduated *cum laude* for my [thesis](#) on Learning Causal Representation in Spatio-Temporal Systems

Master of Science, Artificial Intelligence

University of Amsterdam

📅 Sep 2018 – Aug 2020 📍 Amsterdam, Netherlands

- Courses on AI, including ML, DL, NLP, IR, CV, and RL.
- Thesis on "Categorical Normalizing Flows" ([publication](#)), applied for graph/molecule generation.
- Final GPA: 9.5, cum laude (Dutch grading system).

Bachelor of Engineering, Computer Science

Baden-Wuerttemberg Cooperative State University

📅 Oct 2015 – Sep 2018 📍 Stuttgart, Germany

- Cooperative study program in cooperation with Daimler AG/Mercedes-Benz R&D for autonomous driving.
- Final GPA: 1.0 (German grading system).

AWARDS

- Google-internal Peer Award for finding and resolving severe bug in distributed data loading, affecting hundreds of large-scale research experiments globally.
- Won 4th place in the NeurIPS 2020 challenge "[Hateful Memes](#)" of Facebook AI ([paper](#)).
- Best CS undergraduate student at the Baden-Wuerttemberg Cooperative State University 2018.
- Awarded SchülerUni Scholarship 2012/13 for taking university courses during high school.

SKILLS

DL Frameworks: PyTorch (2018-present, R&T), JAX/Flax (2021-present, R&T), Tensorflow (2017-2023, R), Caffe (2016-2017, R) (R - Research, T - Teaching)

Programming languages: Python (2016-present), HTML (2012-present), Java (2012-2019), C (2015-2017)

Additional skills: SLURM cluster (2018-present), git (2015-present), Triton/Pallas (2024-present)

SELECTED PUBLICATIONS (FULL LIST OF PUBLICATIONS)

- Maximilian Beck*, Korbinian Pöppel*, [Phillip Lippe](#)*, Richard Kurle, Patrick M Blies, Günter Klambauer, Sebastian Böck, Sepp Hochreiter: **xLSTM 7B: A Recurrent LLM for Fast and Efficient Inference**. International Conference on Machine Learning (ICML), 2025 ([link](#))
– *Own contributions*: Developing large-scale training framework, kernel optimizations, research on scalability and proposed core improvements.
- Samuele Papa, Riccardo Valperga, David M Knigge, Miltiadis Kofinas, [Phillip Lippe](#), Jan-jakob Sonke, Efstratios Gavves: **How to Train Neural Field Representations: A Comprehensive Study and Benchmark**. Computer Vision and Pattern Recognition (CVPR), 2024 ([link](#))
– *Own contributions*: Core contributor of library development in JAX, proposed and optimized vmap parallelization, advised in model development.
- Davide Talon, [Phillip Lippe](#), Stuart James, Alessio Del Bue, Sara Magliacane: **Towards the Reusability and Compositionality of Causal Representations**. Third Conference on Causal Learning and Reasoning (CLear), 2024 ([link](#)) [[Oral](#)]
– *Own contributions*: Worked on model development, experiment design, and theoretical results.. Contributed to writing.
- [Phillip Lippe](#), Bastiaan S. Veeling, Paris Perdikaris, Richard E. Turner, Johannes Brandstetter: **PDE-Refiner: Achieving Accurate Long Rollouts with Neural PDE Solvers**. Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS), 2023 ([link](#)). [[Spotlight](#)]
- Sindy Löwe, [Phillip Lippe](#), Francesco Locatello, Max Welling: **Rotating Features for Object Discovery**. Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS), 2023 ([link](#)). [[Oral](#)]
– *Own contributions*: Advised in model development, experiment design and model optimization. Contributed to writing.
- [Phillip Lippe](#), Sara Magliacane, Sindy Löwe, Yuki M. Asano, Taco Cohen, Efstratios Gavves: **BISCUIT: Causal Representation Learning from Binary Interactions**. The 39th Conference on Uncertainty in Artificial Intelligence (UAI), 2023 ([link](#)). [[Spotlight](#)]
- [Phillip Lippe](#), Sara Magliacane, Sindy Löwe, Yuki M. Asano, Taco Cohen, Efstratios Gavves: **Causal Representation Learning for Instantaneous and Temporal Effects in Interactive Systems**. International Conference on Learning Representations (ICLR), 2023 ([link](#)).
- Adeel Pervez, [Phillip Lippe](#), Efstratios Gavves: **Differentiable Mathematical Programming for Object-Centric Representation Learning**. International Conference on Learning Representations (ICLR), 2023 ([link](#)).
– *Own contributions*: Proposed application to object-centric learning. Advised in model development and experimental setups. Contributed to writing.
- Johann Brehmer, Pim de Haan, [Phillip Lippe](#), Taco Cohen: **Weakly supervised causal representation learning**. Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS), 2022 ([link](#)).
– *Own contributions*: Advised in theory and model development, created datasets, implemented causal discovery pipeline, and contributed to writing.
- [Phillip Lippe](#), Sara Magliacane, Sindy Löwe, Yuki M. Asano, Taco Cohen, Efstratios Gavves: **CITRIS: Causal Identifiability from Temporal Intervened Sequences**. International Conference on Machine Learning (ICML), 2022 ([link](#)). [[Spotlight](#)]
- [Phillip Lippe](#), Taco Cohen and Efstratios Gavves: **Efficient Neural Causal Discovery without Acyclicity Constraints**. International Conference on Learning Representations (ICLR), 2022 ([link](#)).
- [Phillip Lippe](#) and Efstratios Gavves: **Categorical Normalizing Flows via Continuous Transformations**. International Conference on Learning Representations (ICLR), 2021 ([link](#)).
- [Phillip Lippe](#), Nithin Holla, Shantanu Chandra, Santhosh Rajamanickam, Georgios Antoniou, Ekaterina Shutova, and Helen Yannakoudakis: **A Multimodal Framework for the Detection of Hateful Memes**. NeurIPS 2020 Competition Track, Facebook AI Hateful Memes ([link](#)).

I have served as reviewer for the following venues: [ECCV-2020](#), [ICCV-2021](#), [CausalUAI-2021](#), [ICLR-2022](#), [CVPR-2022](#), [CLear-2022](#), [NeurIPS-2022](#), [CRL-2022](#), [CML4Impact-2022](#), [CDS-2022](#), [nCSI-2022](#), [ICLR-2023](#), [CLear-2023](#), [TSRL4H-2023](#), [Physics4ML-2023](#), [UAI-2023](#), [NeurIPS-2023](#), [ICML-2024](#). I co-organize the [CORR](#) workshop at CVPR-2024.

SELECTED PUBLIC TALKS (FULL LIST OF TALKS)

- 03/2024 Invited Talk at the [Deep Thinking Hour \(UvA\)](#) on "Training Models at Scale"
- 02/2024 Invited Talk at the [CARE Talk Series by Valencelabs](#) on "BISCUIT: CRL from Binary Interactions"
- 02/2024 Invited Talk at the [Bellairs Workshop on Causality](#) on "On Practical Challenges of Scaling CRL"
- 09/2023 Invited Talk at the [AI4Science Talk Series](#) on "Achieving Accurate Long Rollouts with Neural PDE Solvers"
- 02/2023 Invited Talk at the [Rising Stars in AI Symposium](#) at KAUST on "Causal Representation Learning"
- 12/2022 Invited Talk at the [Google Student Developer Club, University of Augsburg](#) on "ML with JAX and Flax"
- 08/2022 Invited Talk at the [First Workshop on Causal Representation Learning](#) at UAI 2022 on "Learning Causal Variables from Temporal Sequences with Interventions"

SELECTED OPEN-SOURCE CONTRIBUTIONS (VERSION: SEP 14, 2025)

[phlippe/uvadlc_notebooks](#)

★ 2,973 stars 🍴 644 forks

This repository contains a series of Deep Learning tutorials written in PyTorch and JAX (Jupyter notebooks). The topics range from introductions to PyTorch/JAX, Optimization and Initialization, to recent advances in (Vision) Transformers, Generative Models, and Distributed Training at Scale. The tutorials are easily accessible and executable via its [website](#). So far, 40 tutorials have been authored by myself, and 18 notebooks by collaborators.

[Lightning-AI/tutorials](#)

★ 319 stars 🍴 80 forks

This repository contains the official tutorials for PyTorch Lightning. In collaboration with the Lightning-AI team, I have integrated my tutorials from the [phlippe/uvadlc_notebooks](#) repository. This included common steps for large, open-source libraries, such as code reviews, PEP8 formatting, automatic github actions, and similar.