



NeurIPS 2024
Accepted papers

1. Tristan Cinquin, Marvin Pförtner, Vincent Fortuin, **Philipp Hennig**, **Robert Bamler**
FSP-Laplace: Function-Space Priors for the Laplace Approximation in Bayesian Deep Learning
2. Sebastian Damrich, **Philipp Berens**, **Dmitry Kobak**
Persistent Homology for High-dimensional Data Based on Spectral Methods
3. **Vishaal Udandarao**, **Karsten Roth**, Sebastian Dziadzio, Ameya Prabhu, Mehdi Cherti, Oriol Vinyals, Olivier Henaff, Samuel Albanie, Zeynep Akata, **Matthias Bethge**
A Practitioner's Guide to Real-World Continual Multimodal Pretraining
4. Vishaal Udandarao, Ameya Prabhu, Adhiraj Ghosh, Yash Sharma, Philip Torr, Adel Bibi, Samuel Albanie, Matthias Bethge
No "Zero-Shot" Without Exponential Data: Pretraining Concept Frequency Determines Multimodal Model Performance
5. Ori Press, Andreas Hochlehnert, Ameya Prabhu, Vishaal Udandarao, Ofir Press, **Matthias Bethge**
CiteME: Can Language Models Accurately Cite Scientific Claims?
6. Matthias Tangemann, Matthias Kümmerer, **Matthias Bethge**
Object segmentation from common fate: Motion energy processing enables human-like zero-shot generalization to random dot stimuli
7. Ameya Prabhu, Vishaal Udandarao, Philip Torr, **Matthias Bethge**, Adel Bibi, Samuel Albanie
Efficient Lifelong Model Evaluation in an Era of Rapid Progress
8. Anna Mészáros, Patrik Reizinger, Szilvia Ujváry, **Wieland Brendel**, Ferenc Huszar
Rule Extrapolation in Language Modeling: A Study of Compositional Generalization on OOD Prompts
9. Roland S. Zimmermann, David Klindt, **Wieland Brendel**
Measuring Per-Unit Interpretability at Scale Without Humans
10. Tankred Saanum, **Peter Dayan**, Eric Schulz
Simplifying Latent Dynamics with Softly State-Invariant World Models
11. Jack Merullo, **Carsten Eickhoff**, Ellie Pavlick
Talking Heads: Understanding Inter-Layer Communication in Transformer Language Models
12. Haiwen Huang, Songyou Peng, Dan Zhang, **Andreas Geiger**
Renovating Names in Open-Vocabulary Segmentation Benchmarks
13. Shenyan Gao, Jiazhi Yang, Li Chen, Kashyap Chitta, Yihang Qiu, **Andreas Geiger**, Jun Zhang, Hongyang Li
Vista: A Generalizable Driving World Model with High Fidelity and Versatile Controllability
14. Daniel Dauner, Marcel Hallgarten, Tianyu Li, Xinshuo Weng, Zhiyu Huang, Zetong Yang, Hongyang Li, Igor Gilitschenski, Boris Ivanovic, Marco Pavone, **Andreas Geiger**, Kashyap Chitta
NAVSIM: Data-Driven Non-Reactive Autonomous Vehicle Simulation and Benchmarking
15. Alexander Lappe, Anna Bognár, Ghazaleh Ghamkahri Nejad, Albert Mukovskiy, Lucas Martini, **Martin Giese**, Rufin Vogels
Parallel Backpropagation for Shared-Feature Visualization
16. Vivian Nastl, **Moritz Hardt**
Do causal predictors generalize better to new domains?

17. Ricardo Dominguez-Olmedo, **Moritz Hardt**, **Celestine Mender-Dünner**
Questioning the Survey Responses of Large Language Models
18. **Celestine Mender-Dünner**, Gabriele Carovano, **Moritz Hardt**
An engine not a camera: Measuring performative power of online search
19. André F. Cruz, **Celestine Mender-Dünner**, **Moritz Hardt**
Evaluate calibration of language models with folktexts
20. Hrittik Roy, Marco Miani, Carl Henrik Ek, **Philipp Hennig**, Marvin Pförtner, Lukas Tatzel, Søren Hauberg
Reparameterization invariance in approximate Bayesian inference
21. Jonathan Wenger, Kaiwen Wu, **Philipp Hennig**, Jacob Gardner, Geoff Pleiss, John Cunningham
Computation-Aware Gaussian Processes: Model Selection And Linear-Time Inference
22. Irene Huang, Wei Lin, Muhammad Mirza, Jacob Hansen, Sivan Doveh, Victor Butoi, Roei Herzig, Assaf Arbelle, **Hilde Kuehne**, Trevor Darrell, Chuang Gan, Aude Oliva, Rogerio Feris, Leonid Karlinsky
ConMe: Rethinking Evaluation of Compositional Reasoning for Modern VLMs
23. Felix Petersen, Christian Borgelt, Tobias Sutter, **Hilde Kuehne**, Oliver Deussen, Stefano Ermon
Fishers and Hessians of Continuous Relaxations
24. Felix Petersen, **Hilde Kuehne**, Christian Borgelt, Julian Welzel, Stefano Ermon
Convolutional Differentiable Logic Gate Networks
25. Jan-Niklas Dihlmann, Arjun Majumdar, Andreas Engelhardt, Raphael Braun, **Hendrik PA Lensch**
Subsurface Scattering for Gaussian Splatting
26. Robi Bhattacharjee, **Ulrike Luxburg**
Auditing Local Explanations is Hard
27. Julius Vetter, Guy Moss, Cornelius Schröder, Richard Gao, **Jakob H Macke**
Sourcerer: Sample-based Maximum Entropy Source Distribution Estimation
28. Matthijs Pals, A Erdem Sağtekin, Felix Pei, Manuel Gloeckler, **Jakob H Macke**
Inferring stochastic low-rank recurrent neural networks from neural data
29. Jaivardhan Kapoor, Auguste Schulz, Julius Vetter, Felix Pei, Richard Gao, **Jakob H Macke**
Latent Diffusion for Neural Spiking Data
30. Joachim Baumann, **Celestine Mender-Dünner**
Algorithmic Collective Action in Recommender Systems: Promoting Songs by Reordering Playlists
31. István Sáráandi, **Gerard Pons-Moll**
Neural Localization Fields for Continuous 3D Human Pose and Shape Estimation
32. Yuxuan Xue, Xianghui Xie, Riccardo Marin, **Gerard Pons-Moll**
Human 3Diffusion: Realistic Avatar Creation via Explicit 3D Consistent Diffusion Models
33. Ahmad-Reza Ehyaei, Golnoosh Farnadi, **Samira Samadi**
Wasserstein Distributionally Robust Optimization through the Lens of Structural Causal Models and Individual Fairness
34. Mohammad-Amin Charusaie, **Samira Samadi**
A Unifying Post-Processing Framework for Multi-Objective Learn-to-Defer Problems
35. Siyuan Guo, Chi Zhang, Karthika Mohan, Ferenc Huszar, **Bernhard Schölkopf**
Do Finetti: On Causal Effects for Exchangeable Data
36. Carl-Johann Simon-Gabriel, Alessandro Barp, **Bernhard Schölkopf**, Lester Mackey
Metrizing Weak Convergence with Maximum Mean Discrepancies
37. Giorgio Piatti, Zhijing Jin, Max Kleiman-Weiner, **Bernhard Schölkopf**, Mrinmaya Sachan, Rada Mihalcea
Cooperate or Collapse: Emergence of Sustainability in a Society of LLM Agents

38. Robin Chan, Reda Boumasmoud, Anej Svete, Yuxin Ren, Qipeng Guo, Zhijing Jin, Shauli Ravfogel, Mrinmaya Sachan, **Bernhard Schölkopf**, Mennatallah El-Assady, Ryan Cotterell
On Affine Homotopy between Language Encoders
39. Sergio Garrido Mejia, Patrick Blöbaum, **Bernhard Schölkopf**, Dominik Janzing
Causal vs. Anticausal merging of predictors
40. Goutham Rajendran, Simon Buchholz, Bryon Aragam, **Bernhard Schölkopf**, Pradeep Ravikumar
From Causal to Concept-Based Representation Learning
41. Jonathan Thomm, Aleksandar Terzic, Giacomo Camposampiero, Michael Hersche, **Bernhard Schölkopf**, Abbas Rahimi
Limits of Transformer Language Models on Learning to Compose Algorithms