

1. Ilze Amanda Auzina, Joschka Strüber, Sergio Hernández-Gutiérrez, Shashwat Goel, Ameya Pandurang Prabhu, **Matthias Bethge**
Intrinsic Credit Assignment for Long Horizon Interaction
2. Jana Zeller, Thaddäus Wiedemer, Fanfei Li, Thomas Klein, Prasanna Mayilvahanan, **Matthias Bethge, Felix Wichmann, Ryan Cotterell, Wieland Brendel**
MentisOculi: Revealing the Limits of Reasoning with Mental Imagery
3. Ben Rank, Hardik Bhatnagar, Ameya Pandurang Prabhu, Shira Eisenberg, Karina Nguyen, **Matthias Bethge, Maksym Andriushchenko**
PostTrainBench: Can LLM Agents Automate LLM Post-Training?
4. Cagatay Yildiz, Nishaanth Kanna, Nitin Sharma, **Matthias Bethge, Beyza Ermis**
Investigating Continual Pretraining in Large Language Models: Insights and Implications
5. Shruti Joshi, Aaron Mueller, David Klindt, **Wieland Brendel, Dhanya Sridhar, Patrik Reizinger**
Position: Causality is Key for Interpretability Claims to Generalise
6. Jack Brady, **Bernhard Schölkopf, Thomas Kipf, Simon Buchholz, Wieland Brendel**
Generation is Required for Data-Efficient Perception
7. Manuel Traub, **Martin V Butz**
Looking Locally: Object-Centric Vision Transformers as Foundation Models for Efficient Segmentation
8. Moritz Brösamle, **Stephan Eckstein**
The Expressive Power of Low Precision Softmax Transformers with (Summarized) Chain-of-Thought
9. Fabian Morelli, **Stephan Eckstein**
Partial Fusion of Neural Networks: Efficient Tradeoffs Between Ensembles and Weight Aggregation
10. Haofei Xu, Rundi Wu, Philipp Henzler, Nikolai Kalischek, Michael Oechsle, Fabian Manhardt, Marc Pollefeys, **Andreas Geiger, Federico Tombari, Michael Niemeyer**
PointDiT: Pixel-Space Diffusion for Monocular Geometry Estimation
11. Ricardo Dominguez-Olmedo, **Bernhard Schölkopf, Moritz Hardt**
Computational Arbitrage in AI Model Markets
12. Mina Remeli, **Moritz Hardt**
Correct looks better: Pairwise comparisons reveal accuracy rankings
13. Yatong Chen, Guanhua Zhang, **Moritz Hardt**
Leaderboard Incentives: Model Rankings under Strategic Post-Training
14. Nikhil Chandak, Shashwat Goel, Ameya Pandurang Prabhu, **Moritz Hardt, Jonas Geiping**
Curating the Future: A Scalable Recipe for Training Open-Ended Forecasters
15. Tim Weiland, **Philipp Hennig**
Scalable Bayesian Inference for Nonlinear Conservation Laws

16. Arik Reuter, Anish Dhir, Cristiana Diaconu, Jake Robertson, Ole Ossen, **Frank Hutter**,
Adrian Weller, Mark van der Wilk, **Bernhard Schölkopf**
Use What You Know: Causal Foundation Models with Partial Graphs
17. Herilalaina Rakotoarison, Steven Adriaensen, Tom Viering, Samuel Gabriel Müller,
Carl Hvarfner, **Frank Hutter**, Eytan Bakshy
 α -PFN: Fast Entropy Search via In-Context Learning
18. Tom Sühr, Florian Dorner, Olawale Salaudeen, **Augustin Kelava**, **Samira Samadi**
Position: Stop evaluating AI with human tests, develop principled, AI-specific tests instead
19. Han Wang, Weijie Wang, Jiaqi Liu, **Hilde Kuehne**, Nicu Sebe
Deep Trajectory Supervision: Deep Supervision Strikes Back
20. Stefan Wahl, Raphaela Schenk, Ali Farnoud, **Jakob Macke**, Daniel Gedon
A Probabilistic Framework for LLM-Based Model Discovery
21. Manuel Glöckler, Jose Pedro JP Manzano-Patron, Stamatios Sotiropoulos, Cornelius
Schröder, **Jakob Macke**
Scalable Simulation-Based Model Inference with Test-Time Complexity Control
22. Mikel Zhobro, Andreas René Geist, **Georg Martius**
Learning 3D-Gaussian Simulators from RGB Videos
23. Anselm Paulus, Andreas René Geist, Vit Musil, Sebastian Hoffmann, **Georg Martius**
SoftJAX & SoftTorch: Empowering Automatic Differentiation Libraries with Informative Gradients
24. Marco Bagatella, Mert Albaba, Jonas Hübötter, **Georg Martius**, Andreas Krause
Test-time Offline Reinforcement Learning on Goal-related Experience
25. **Alina Wernick**, **Kristof Meding**
Position: EU AI Act's Research Exemptions Can Break the Publication Norms of Major AI Conferences
26. Ruta Binkyte, Ivaxi Sheth, Zhijing Jin, Mohammad Havaei, **Bernhard Schölkopf**,
Mario Fritz
Position: Trustworthy AI Suffers from Invariance Conflicts and Causality is The Solution
27. Sirui Lu, Zhijing Jin, Terry Zhang, Pavel Kos, Juan Cirac, **Bernhard Schölkopf**
Position: LLM for Physics Research Requires Domain-Specialized Training and Tooling
28. Samuel Simko, Punya Pandey, Zhijing Jin, **Bernhard Schölkopf**
Training with Honeypots: Reshaping How LLMs Fail
29. Sawal Acharya, Terry Zhang, Andrew Kim, Anahita Haghighat, Xianlin Sun, Pepijn
Cobben, Rahul Shrestha, Maximilian Mordig, Jacob Emmerson, Furkan Danisman,
Yuen Chen, Clijo Jose, Andrei Muresanu, Justin Cui, Jiarui Liu, Yahang Qi, Punya
Pandey, Yinya Huang, **Bernhard Schölkopf**, Zhijing Jin
CauSciBench: Evaluating LLM Causal Inference for Scientific Research
30. Ziheng Chen, Xiaojun Wu, **Bernhard Schölkopf**, Nicu Sebe
Riemannian Networks over Full-Rank Correlation Matrices
31. David Guzman Piedrahita, Dave Banerjee, Changling Li, Terry Zhang, Kevin Blin,
Samuel Simko, Punya Pandey, Irene Strauss, Rada Mihalcea, **Bernhard Schölkopf**,
Zhijing Jin
Position: Safe Models Do Not Guarantee Safe Societies: The Case for Sociopolitical Risk
32. Onno Eberhard, **Claire Vernade**, Michael Muehlebach
Commit to the Bit: Reactive Reinforcement Learning Done Right