

## **NeurIPS 2023 – Accepted papers**

1. **Leonard Salewski**, Isabel Rio-Torto, Stephan Alaniz, **Eric Schulz**, **Zeynep Akata**  
*In-Context Impersonation Reveals Large Language Models' Strengths and Biases*
2. Julian Coda-Forno, Marcel Binz, **Zeynep Akata**, Matt Botvinick, Jane Wang, **Eric Schulz**  
*Meta-in-context learning in large language models*
3. ABHRA CHAUDHURI, **Massimiliano Mancini**, **Zeynep Akata**, Anjan Dutta  
*Transitivity Recovering Decompositions: Interpretable and Robust Fine-Grained Relationships*
4. Ilze Amanda Auzina, **Çağatay Yıldız**, Sara Magliacane, **Matthias Bethge**, Efstratios Gavves  
*Modulated Neural ODEs*
5. Thaddäus Wiedemer, Prasanna Mayilvahanan, **Matthias Bethge**, **Wieland Brendel**  
*Compositional Generalization from First Principles*
6. Ori Press, Steffen Schneider, Matthias Kümmerer, **Matthias Bethge**  
*RDumb: A simple approach that questions our progress in continual test-time adaptation*
7. Roland S. Zimmermann, Thomas Klein, **Wieland Brendel**  
*Scale Alone Does not Improve Mechanistic Interpretability in Vision Models*
8. **Tankred Saanum**, Noemi Elteto, **Peter Dayan**, Marcel Binz, **Eric Schulz**  
*Reinforcement Learning with Simple Sequence Priors*
9. Maximilian Mueller, Tiffany Vlaar, David Rolnick, **Matthias Hein**  
*Normalization Layers Are All That Sharpness-Aware Minimization Needs*
10. Naman Deep Singh, Francesco Croce, **Matthias Hein**  
*Revisiting Adversarial Training for ImageNet: Architectures, Training and Generalization across Threat Models*
11. **Jonathan Schmidt**, **Philipp Hennig**, Jörg Nick, Filip Tronarp  
*The Rank-Reduced Kalman Filter: Approximate Dynamical-Low-Rank Filtering In High Dimensions*
12. Nathanael Bosch, **Philipp Hennig**, Filip Tronarp  
*Probabilistic Exponential Integrators*
13. Runa Eschenhagen, Alexander Immer, Richard Turner, Frank Schneider, **Philipp Hennig**  
*Kronecker-Factored Approximate Curvature for Modern Neural Network Architectures*
14. Agustinus Kristiadi, Felix Dangel, **Philipp Hennig**  
*The Geometry of Neural Nets' Parameter Spaces Under Reparametrization*
15. **Michael Kirchhof**, Bálint Mucsányi, Seong Joon Oh, **Dr. Enkelejda Kasneci**  
*URL: A Representation Learning Benchmark for Transferable Uncertainty Estimates*

16. **Moritz Haas**, David Holzmüller, **Ulrike Luxburg**, Ingo Steinwart  
*Mind the spikes: Benign overfitting of kernels and neural networks in fixed dimension*
17. Basile Confavreux, Poornima Ramesh, Pedro Goncalves, **Jakob H Macke**, Tim Vogels  
*Meta-learning families of plasticity rules in recurrent spiking networks using simulation-based inference*
18. Jonas Wildberger, Maximilian Dax, Simon Buchholz, Stephen Green, **Jakob H Macke**, **Bernhard Schölkopf**  
*Flow Matching for Scalable Simulation-Based Inference*
19. Richard Gao, Michael Deistler, **Jakob H Macke**  
*Generalized Bayesian Inference for Scientific Simulators via Amortized Cost Estimation*
20. Marco Bagatella, **Georg Martius**  
*Goal-conditioned Offline Planning from Curious Exploration*
21. Cansu Sancaktar, Justus Piater, **Georg Martius**  
*Regularity as Intrinsic Reward for Free Play*
22. Pavel Kolev, **Georg Martius**, Michael Muehlebach  
*Online Learning under Adversarial Nonlinear Constraints*
23. Andrii Zadaianchuk, Maximilian Seitzer, **Georg Martius**  
*Object-Centric Learning for Real-World Videos by Predicting Temporal Feature Similarities*
24. Simon Buchholz, Goutham Rajendran, Elan Rosenfeld, Bryon Aragam, **Bernhard Schölkopf**, Pradeep Ravikumar  
*Learning Linear Causal Representations from Interventions under General Nonlinear Mixing*
25. Zeju Qiu, Weiyang Liu, Haiwen Feng, Yuxuan Xue, Yao Feng, Zhen Liu, Dan Zhang, Adrian Weller, **Bernhard Schölkopf**  
*Controlling Text-to-Image Diffusion by Orthogonal Finetuning*
26. Laurence Midgley, Vincent Stimper, Javier Antorán, Emile Mathieu, **Bernhard Schölkopf**, José Miguel Hernández-Lobato  
*SE(3) Equivariant Augmented Coupling Flows*
27. Siyuan Guo, Viktor Toth, **Bernhard Schölkopf**, Ferenc Huszar  
*Causal de Finetti: On the Identification of Invariant Causal Structure in Exchangeable Data*
28. Liang Wendong, Armin Kekić, Julius von Kügelgen, Simon Buchholz, Michel Besserve, Luigi Gresele, **Bernhard Schölkopf**  
*Causal Component Analysis*
29. Zhijing Jin, Yuen Chen, Felix Leeb, Luigi Gresele, Ojasv Kamal, Zhiheng LYU, Kevin Blin, Fernando Gonzalez Adauto, Max Kleiman-Weiner, Mrinmaya Sachan, **Bernhard Schölkopf**  
*CLadder: Assessing Causal Reasoning in Language Models*

30. Julius von Kügelgen, Michel Besserve, Liang Wendong, Luigi Gresele, Armin Kekić, Elias Bareinboim, David Blei, **Bernhard Schölkopf**  
*Nonparametric Identifiability of Causal Representations from Unknown Interventions*
31. Marco Fumero, Florian Wenzel, Luca Zancato, Alessandro Achille, Emanuele Rodolà, Stefano Soatto, Bernhard Schölkopf, Francesco Locatello  
*Leveraging sparse and shared feature activations for disentangled representation learning*
32. Junhyung Park, Simon Buchholz, **Bernhard Schölkopf**, Krikamol Muandet  
*A Measure-Theoretic Axiomatisation of Causality*
33. Cian Eastwood, Shashank Singh, Andrei L Nicolicioiu, Marin Vlastelica Pogančić, Julius von Kügelgen, **Bernhard Schölkopf**  
*Spuriousity Didn't Kill the Classifier: Using Invariant Predictions to Harness Spurious Features*
34. Adrián Javaloy, **Pablo Sanchez-Martin, Isabel Valera**  
*Causal normalizing flows: from theory to practice*
35. Alexandre Marthe, Aurélien Garivier, **Claire Vernade**  
*Beyond Average Reward in Markov Decision Processes*