machine learning
new perspectives
for science

NeurIPS Papers 2026
Accepted Papers

1. *Reaching a Consensus in Predictive Loops*
   Jiduan Wu, **Rediet Abebe,** Celestine Mendler-Dünner
2. *MedKIT: Evaluating Knowledge Integration and Generalization in Large Language Models*
   Lukas Thede, Yash Kumar, David Chen, Danielle Bitterman, **Matthias Bethge**, Tom Hartvigsen, Zeynep Akata
3. *SignalBench: Comparing Dense Feedback Methods for Long-Horizon Agents*
   Sergio Hernández-Gutiérrez, Matteo Merler, Ilze Amanda Auzina, Joschka Strüber, Ameya Prabhu, **Matthias Bethge**
4. *DataComp-VLM: Improved open datasets for Vision-Language Models*
   Matteo Farina, Vishaal Udandarao, Thao Nguyen, Selim Kuzucu, Maximilian Böther, Andreas Hochlehnert, Adhiraj Ghosh, Marianna Nezhurina, Karsten Roth, Joschka Strüber, Yuhui Zhang, Sebastian Dziadzio, Elaine Sui, Soumya S Jahagirdar, Dhruba Ghosh, Hasan Hammoud, Thomas De Min, Simone Caldarella, Muhammad Jehanzeb Mirza, Sedrick Keh, Mehdi Cherti, **Hilde Kuehne**, Bernt Schiele, Serena Yeung-Levy, Muhammad Ferjad Naeem, Federico Tombari, Ana Klimovic, Elisa Ricci, **Matthias Bethge**, Sewoong Oh, Ameya Prabhu, Alessio Tonioni, Jenia Jitsev, Massimiliano Mancini, Ludwig Schmidt, Nikhil Parthasarathy
5. *Evaluating Neural Data Tokenizers: A Framework for Assessing Learned Representations of Spiking Activity*
   Federico D'Agostino, Alex Gilbert, Susanne Keller, Jaivardhan Kapoor, Nicolas Reategui, Vedang Lad, Matthias Kümmerer, Jai Bhagat, Emin Orhan, Hasan A Bedel, Nikos Karantzas, Fabian Sinz, Surya Ganguli, **Jakob H Macke**, Sophia Sanborn, Alexander Ecker, Katrin Franke, **Matthias Bethge**, Andreas Tolias, Konstantin Willeke
6. *LAION-BVD: A 10-Million-Hour Open Video Dataset for Multimodal Pre-training*
   Andreas Hochlehnert, Marianna Nezhurina, Mehdi Cherti, Andrej Radonjic, Thaddäus Wiedemer, Christoph Schuhmann, Romain Beaumont, **Wieland Brendel**, **Bernhard Schölkopf**, A. Sophia Koepke, Jenia Jitsev, **Matthias Bethge**
7. *RevengeBench: Reverse Engineering Code-Space Policies from Behavioral Experiments*
   Babak Rahmani, Sebastian Dziadzio, Joschka Strüber, Sergio Hernández-Gutiérrez, **Matthias Bethge**
8. *Stabilizing Extrapolation in Looped Transformers via Learned Stochastic Stopping*
   Hsun-Yu Kuo, El Mahdi Chayti, Patrik Reizinger, **Wieland Brendel**, Martin Jaggi
9. *LittleLearner: Language Models Under Pedagogically-Controlled Knowledge Exposure*
   Fanfei Li, Jana Zeller, Manuel Prada-Corral, Thaddäus Wiedemer, Prasanna Mayilvahanan, Ryan Cotterell, **Wieland Brendel**

10. *Unifying Goal-Conditioned RL and Unsupervised Skill Learning via Control-Maximization*
    Alireza Modirshanechi, Benjamin Eysenbach, **Peter Dayan**, Eric Schulz
11. *Fast Wasserstein rates for estimating probability distributions of probabilistic graphical models*
    Daniel Bartl, **Stephan Eckstein**
12. *From Topology to Retrieval: Decoding Embedding Spaces with Unified Signatures*
    Florian Rottach, William Rudman, Bastian Rieck, **Harrisen Scells**, **Carsten Eickhoff**
13. *Understanding Multi-View Transformers*
    Julien Gaubil, Michal Stary, Louis Martinez, **Andreas Geiger**, Maks Ovsjanikov, Ayush Tewari, Vincent Sitzmann
14. *FutureSim: Replaying World Events to Evaluate Adaptive Agent*s
    Shashwat Goel, Nikhil Chandak, Arvindh Arun, Ameya Prabhu, Steffen Staab, **Moritz Hardt**, Maksym Andriushchenko, Jonas Geiping
15. *Unlearning That Lasts: Utility-Preserving, Robust, and Almost Irreversible Forgetting in LLMs*
    Naman Deep Singh, Maximilian Mueller, Amit Peleg, Francesco Croce, **Matthias Hein**
16. *Learning POMDP World Models from Observations with Language-Model Priors*
    Valentin Six, Frederik Panse, Mathis Fajeau, Lancelot Da Costa, Mridul Sharma, Alfonso Amayuelas, Tim Xiao, David Hyland, **Philipp Hennig**, **Bernhard Schölkopf**
17. *Conditioning Gaussian Processes on Almost Anything*
    Henry Moss, Lachlan Astfalck, Tom Cowperthwaite, Colin Doumont, Samuel Willis, Christopher Nemeth, **Philipp Hennig**, Andrew Zammit-Mangion
18. *Closed-Form Last Layer Optimizatio*n
    Alexandre Galashov, Nathaël Da Costa, Liyuan Xu, **Philipp Hennig**, Arthur Gretton
19. *BO-Arena: An Evolving Benchmark for High-Dimensional Bayesian Optimisatio*n
    Thomas Christie, Colin Doumont, Henry Moss, **Philipp Hennig**
20. *Beyond IID: How General Are Tabular Foundation Models, Really?*
    Lennart Purucker, Andrej Tschalzev, Nick Erickson, David Holzmüller, Alan Arazi, Alexander Pfefferle, Mustafa Tajjar, Gioia Blayer, Gael Varoquaux, **Frank Hutter**
21. *Understanding the Surprising Generalization Properties of Tabular Foundation Models*
    Nour Shaheen, Junwei (Jeremy) Ma, Alex Labach, **Frank Hutter**, Anthony L Caterini, Valentin Thomas
22. *Encoding RNA Topology into Synthetic Alignments for 3D Structure Prediction*
    Jörg Franke, Iris K Tennie, Michael Uhl, Ryan Koksal, Dominika Matus, Dominik Scheuer, **Frank Hutter**, Rolf Backofen, Frederic Runge
23. *STRABLE: Benchmarking Tabular Machine Learning with Strings*
    Gioia Blayer, Myung Jun Kim, Félix Lefebvre, Lennart Purucker, Alan Arazi, Eilam Shapira, Roi Reichart, **Frank Hutter**, Marine Le Morvan, David Holzmüller, Gael Varoquaux

24. *When is Warmstarting Effective for Scaling Language Models?*
    Neeratyoy Mallik, Maciej Janowski, Johannes Hog, Herilalaina Rakotoarison, Josif Grabocka, **Frank Hutter**, Aaron Klein
25. *MulTaBench: Benchmarking Multimodal Tabular Learning with Text and Image*
    Alan Arazi, Eilam Shapira, Shoham Grunblat, Mor Ventura, Elad Hoffer, Gioia Blayer, David Holzmüller, Lennart Purucker, Gael Varoquaux, **Frank Hutter**, Roi Reichart
26. *TempoPFN: Synthetic Pre-training of Linear RNNs for Zero-shot Time Series Forecasting*
    Vladyslav Moroshan, Julien Siems, Arber Zela, Timur Carstensen, **Frank Hutter**
27. *AI Construct Lexis: An Ontology of the Hidden Assumptions in AI Evaluation*
    Olawale Salaudeen, Florian E. Dorner, Tom Sühr, Sang Truong, Haoran Zhang, Patrick Bissett, James Fox, Marzyeh Ghassemi, Tom Hartvigsen, Peter Hase, **Augustin Kelava**, Suhas Mahesh, Margaret Mitchell, Jamelle Watson-Daniels, Sanmi Koyejo, Angelina Wang
28. *FalconPerception-HD: High Density Perception via Reinforcement Learning*
    Sofian Chaybouti, YASSER ABDELAZIZ DAHOU DJILALI, Ngoc D Huynh, REDA ALAMI, **Hilde Kuehne**
29. *When Trackers Fail: VLM-Guided Verification and Recovery for Robust Video Object Segmentation*
    Valay Mahesh Bundele, Susmit Agrawal, Mehran Hosseinzadeh, The Nam Nguyen, **Hendrik PA Lensch**
30. *Scaling Laws and Tradeoffs in Recurrent Networks of Expressive Neurons*
    Aaron Spieler, **Georg Martius**, **Anna Levina**
31. *Soft Forward-Backward Representations for Zero-shot Reinforcement Learning with General Utilities*
    Marco Bagatella, Thomas Rupf, **Georg Martius**, Andreas Krause
32. *Drifting Fields are not Conservative*
    Leonard T. Franz, Sebastian Hoffmann, Tim Weiland, **Bernhard Schölkopf**, **Georg Martius**
33. *SAMoR: Motion Modelling for Articulated Objects of Any Skeleton and Topology*
    Yuhao Zhang, **Gerard Pons-Moll**, Tolga Birdal
34. *CoTs as Probabilistic Programs: A Programmatic View of Thinking Step-by-Step in Language Models*
    Kyle Richardson, Yu Feng, Poorva Garg, Junyan Cheng, Dan Roth, Guy Van den Broeck
35. *Operads for compositional reasoning in LLMs*
    Nathaniel Bottman, Yinhong Liu, **Kyle Richardson**
36. *Controlling for Omitted Variable Bias in Deep Neural Networks*
    Manuel Pfeuffer, Roshan P Rane, **Kerstin Ritter**, Sonja Greven
37. *Redefining Instance Matching: A Unified Framework for Part-Aware Matching in Panoptic Segmentation Evaluation*
    Erik Großkopf, Soumya Snigdha Kundu, Hendrik Möller, Nicolas Münster, Mehdi Astaraki, Paula T Buzduga, **Kerstin Ritter**, Benedikt Wiestler, Jan Kirschke, Jonathan Shapey, Tom Vercauteren, Florian Kofler

38. *Flow Matching with In-Context Priors for Out-of-Distribution Brain Dynamics*
    Sam Gijsen, Michał Łukomski, Marc-Andre Schulz, **Kerstin Ritter**
39. *Foveated BagNet: Inherent Interpretability Does Not Exclude Global Context*
    Holger Heidrich, Sarah Müller, **Andreas Schilling**
40. *PENEX: AdaBoost-Inspired Neural Network Regularization*
    Klaus-Rudolf Kladny, **Bernhard Schölkopf**, Michael Muehlebach
41. *PVFormer: Proper Velocity Transformer for Stable and Scalable Hyperbolic Representation Learning*
    Xianglong Shi, Nicu Sebe, **Bernhard Schölkopf**, Ziheng Chen
42. *Deriving Hyperparameter Scaling Laws via Modern Optimization Theory*
    Egor Shulgin, Jörg Franke, Dimitri von Rütte, Tianyue Zhang, Niccolò Ajroldi, Korbinian Pöppel, **Bernhard Schölkopf**, Aaron Klein, Peter Richtarik, Antonio Orvieto
43. *SuperSycophantic: Stress-Testing Frontier LLMs from Single- to Multi-Turn Sycophancy*
    Terry J. C. Zhang, Oscar S Yasunaga, Wenyuan Jiang, Jessica Bo, Florent Draye, Aydin Javadov, Arth Singh, Lucille Muir, VICTORIA OLDEMBURGO DE MELLO, **Bernhard Schölkopf**, Zhijing Jin
44. *AM-Bench: A Unified Taxonomy and Evaluation Suite for Agentic Misalignment*
    Eric Zhang, Terry J. C. Zhang, Chijioke Ugwuanyi, Jerick Shi, **Bernhard Schölkopf**, Zhijing Jin
45. *GT-HarmBench: Benchmarking AI Safety Risks Through the Lens of Game Theory*
    Pepijn Cobben, Xuanqiang A Huang, Thao Pham, Isabel Dahlgren, **Bernhard Schölkopf**, Terry J. C. Zhang, Zhijing Jin
46. *DiffRatio: Training One-Step Diffusion Models Without Teacher Supervision*
    Wenlin Chen, Mingtian Zhang, Jiajun He, Zijing Ou, José Miguel Hernández-Lobato, **Bernhard Schölkopf**, David Barber
47. *Perception for Action in Latent World Models*
    Petr Ivashkov, Randall Balestriero, **Bernhard Schölkopf**
48. *The Alien Space of Science: Sampling Coherent but Cognitively Unavailable Research Directions*
    Alejandro H. Artiles, Martin Weiss, Levin Brinkmann, Iyad Rahwan, **Bernhard Schölkopf**, Chris Pal, Hugo Larochelle, Anirudh Goyal, Nasim Rahaman
49. *Causality can systematically address the monsters under the bench(marks)*
    Felix Leeb, Zhijing Jin, **Bernhard Schölkopf**
50. *Umbilic Multinomial Logistic Regression*
    Zihan Su, Nicu Sebe, **Bernhard Schölkopf**, Ziheng Chen
51. *Intrinsically Interpretable Attention via Sparse Post-Training*
    Florent Draye, Anson Lei, Hsiao-Ru Pan, Ingmar Posner, **Bernhard Schölkopf**
52. *Transferability for General Reasoning: An Automated Curriculum for Multi-Domain LLM RL*
    Yongjin Yang, Jiarui Liu, Yinghui He, Lechen Zhang, **Bernhard Schölkopf**, Zhijing Jin
53. *Corruptions of Supervised Learning Problems: Typology and Mitigations*
    Laura Iacovissi, Nan Lu, **Robert Williamson**