


**Poster Session 1 -
10.45 a.m. (CET)**
#13 *Evaluating Representation Learning on the Protein Structure Universe*

Arian Rokkum Jamasb, Alex Morehead, Zuobai Zhang, Chaitanya K. Joshi, Kieran Didi, Simon V Mathis, Charles Harris, Jian Tang, Jianlin Cheng, Pietro Lio, Tom Leon Blundell

#49 *Synaptic Weight Distributions Depend on the Geometry of Plasticity*

Roman Pogodin, Jonathan Cornford, Arna Ghosh, Gauthier Gidel, Guillaume Lajoie, Blake Aaron Richards

#133 *The Cost of Scaling Down Large Language Models: Reducing Model Size Affects Memory before In-context Learning*

Tian Jin, Nolan Clement, X. Dong, Vaishnavh Nagarajan, Michael Carbin, Jonathan Ragan-Kelley, Gintare Karolina Dziugaite

#266 *LOQA: Learning with Opponent Q-Learning Awareness*

Milad Aghajohari, Juan Agustin Duque, Tim Cooijmans, Aaron Courville

**Orals -
3.45 p.m. (CET)**
Oral 2A
Understanding In-Context Learning in Transformers and LLMs by Learning to Learn Discrete Functions

Satwik Bhattamishra, Arkil Patel, Phil Blunsom, Varun Kanade

Oral 2B
Ghost on the Shell: An Expressive Representation of General 3D Shapes

Zhen Liu, Yao Feng, Yuliang Xiu, Weiyang Liu, Liam Paull, Michael J. Black, Bernhard Scholkopf

Oral 2C
Würstchen: An Efficient Architecture for Large-Scale Text-to-Image Diffusion Models

Pablo Pernias, Dominic Rampas, Mats L. Richter, Christopher Joseph Pal, Marc Aubreville

**Poster Session 2 -
4.30 p.m. (CET)**
#72 *On the Stability of Iterative Retraining of Generative Models on their own Data*

Quentin Bertrand, Avishek Joey Bose, Alexandre Duplessis, Marco Jiralerspong, Gauthier Gidel

#79 *Würstchen: An Efficient Architecture for Large-Scale Text-to-Image Diffusion Models*

Pablo Pernias, Dominic Rampas, Mats L. Richter, Christopher Joseph Pal, Marc Aubreville

#157 *Course Correcting Koopman Representations*

Mahan Fathi, Clement Gehring, Jonathan Pilault, David Kanaa, Pierre-Luc Bacon, Ross Goroshin

#160 *Understanding In-Context Learning in Transformers and LLMs by Learning to Learn Discrete Functions*

Satwik Bhattamishra, Arkil Patel, Phil Blunsom, Varun Kanade

#217 *Object-centric architectures enable efficient causal representation learning*

Amin Mansouri, Jason Hartford, Yan Zhang, Yoshua Bengio

#279 *Ghost on the Shell: An Expressive Representation of General 3D Shapes*

Zhen Liu, Yao Feng, Yuliang Xiu, Weiyang Liu, Liam Paull, Michael J. Black, Bernhard Scholkopf

#284 *Poly-View Contrastive Learning*

Amitis Shidani, Dan Busbridge, Devon Hjelm, Jason Ramapuram, Eeshan Gunesh Dhekane, Russell Webb

#290 *On Diffusion Modeling for Anomaly Detection*

Victor Livernoche, Vineet Jain, Yashar Hezaveh, Siamak Ravanbakhsh

Poster Session 3 - 10.45 a.m. (CET)

#5 *Towards Foundational Models for Molecular Learning on Large-Scale Multi-Task Datasets*

Dominique Beaini, Shenyang Huang, Joao Alex Cunha, Zhiyi Li, Gabriela Moises-cu-Pareja, Oleksandr Dymov, S. Maddrell-Mander, Callum McLean, Frederick Wenkel, Luis Müller, Jama Hussein Mohamud, Alipanah Parviz, Michael Craig, Michal Koziarski, Jiarui Lu, Zhaocheng Zhu, Cristian Gabellini, Kerstin Klaser, Josef Dean, Cas Wognum, Maciej Sypetkowski, Guillaume Rabusseau, Reihaneh Rabbany, Jian Tang, Ioannis Koutis, Christopher Morris, Mirco Ravanelli, Guy Wolf, Prudencio Tossou, Hadrien Mary, Thérance Bois, A. Fitzgibbon, Blazej Banaszewski, Chad Martin, Dominic Masters

#7 *Searching for High-Value Molecules Using Reinforcement Learning and Transformers*

Raj Ghugare, Santiago Miret, Adriana Hugessen, Mariano Phielipp, Glen Berseth

#41 *Tree Cross Attention*

Leo Feng, Frederick Tung, Hossein Hajimirsadeghi, Yoshua Bengio, M. O. Ahmed

#53 *Sufficient conditions for offline reactivation in recurrent neural networks*

Nanda H Krishna, Colin Bredenberg, Daniel Levenstein, Blake Aaron Richards, Guillaume Lajoie

#60 *Ensemble Distillation for Unsupervised Constituency Parsing*

Behzad Shayegh, Yanshuai Cao, Xiaodan Zhu, Jackie CK Cheung, Lili Mou

#98 *GraphPulse: Topological representations for temporal graph property prediction*

Kiarash Shamsi, Farimah Poursafaei, Shenyang Huang, Bao Tran Gia Ngo, Baris Coskunuzer, Cuneyt Gurcan Akcora

#138 *Towards Foundation Models for Knowledge Graph Reasoning*

Mikhail Galkin, Xinyu Yuan, Hesham Mostafa, Jian Tang, Zhaocheng Zhu

#177 *Local Search GFlowNets*

Minsu Kim, Taeyoung Yun, Emmanuel Bengio, Dinghui Zhang, Yoshua Bengio, Sungsoo Ahn, Jinkyoo Park

#181 *Improving Intrinsic Exploration by Creating Stationary Objectives*

Roger Creus Castanyer, Joshua Romoff, Glen Berseth

#214 *Diffusion Generative Flow Samplers: Improving learning signals through partial trajectory optimization*

Dinghui Zhang, Ricky T. Q. Chen, Cheng-Hao Liu, Aaron Courville, Yoshua Bengio

#265 *GOAt: Explaining Graph Neural Networks via Graph Output Attribution*

Shengyao Lu, Keith G Mills, Jiao He, Bang Liu, Di Niu

#307 *Improving Natural Language Understanding with Computation-Efficient Retrieval Augmentation*

Shangyu Wu, Ying Xiong, Yufei CUI, Xue Liu, Buzhou Tang, Tei-Wei Kuo, Chun Jason Xue

Orals - 3.45 p.m. (CET)

Oral 4D

Amortizing intractable inference in large language models
Edward J. Hu, Moksh Jain, Eric Elmoznino, Younesse Kaddar, Guillaume Lajoie, Yoshua Bengio, Nikolay Malkin

Poster Session 4 -
4.30 p.m. (CET)#60 *How connectivity structure shapes rich and lazy learning in neural circuits*

Yuhan Helena Liu, Aristide Baratin, Jonathan Cornford, Stefan Mihalas, Eric Shea-Brown, Guillaume Lajoie

#71 *Pre-Training and Fine-Tuning Generative Flow Networks*

L. Pan, Moksh Jain, Kanika Madan, Yoshua Bengio

#93 *Amortizing intractable inference in large language models*

Edward J. Hu, Moksh Jain, Eric Elmoznino, Younesse Kaddar, Guillaume Lajoie, Yoshua Bengio, Nikolay Malkin

#192 *Empirical Analysis of Model Selection for Heterogeneous Causal Effect Estimation*

Divyat Mahajan, Ioannis Mitliagkas, Brady Neal, Vasilis Syrgkanis

#198 *Closing the Gap between TD Learning and Supervised Learning -- A Generalisation Point of View.*

Raj Ghugare, Matthieu Geist, Glen Berseth, Benjamin Eysenbach

#221 *Balancing Act: Sparse Models with Constrained Disparate Impact*

Meraj Hashemizadeh, Juan Ramirez, Rohan Sukumaran, Golnoosh Farnadi, Simon Lacoste-Julien, Jose Gallego-Posada

#228 *Mechanistically analyzing the effects of fine-tuning on procedurally defined tasks.*

Samyak Jain, Robert Kirk, E. S. Lubana, Robert P. Dick, Hidenori Tanaka, Tim Rocktäschel, Edward Grefenstette, David Krueger

Poster Session 5 - 10.45 a.m. (CET)

#73 *SE(3)-Stochastic Flow Matching for Protein Backbone Generation*

Avishek Joey Bose*, Tara Akhound-Sadegh*, Guillaume Hugué, Kilian FATRAS, Jarrid Rector-Brooks, Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael M. Bronstein, Alexander Tong

#125 *Cycle Consistency Driven Object Discovery*

Aniket Didolkar, Anirudh Goyal, Yoshua Bengio

#145 *TACTiS-2: Better, Faster, Simpler Attentional Copulas for Multivariate Time Series*

Arjun Ashok, Étienne Marcotte, Valentina Zantedeschi, Nicolas Chapados, Alexandre Drouin

#171 *Learning Multi-Agent Communication with Contrastive Learning*

Yat Long Lo, Biswa Sengupta, Jakob Foerster, Michael Noukhovitch

Poster Session 6 - 4.30 p.m. (CET)

#77 *Leveraging Unpaired Data for Vision-Language Generative Models via Cycle Consistency*

Tianhong Li, Sangnie Bhardwaj, Yonglong Tian, Han Zhang, Jarred Barber, Dina Katabi, Guillaume Lajoie, Huiwen Chang, Dilip Krishnan

#142 *Efficient Dynamics Modeling in Interactive Environments with Koopman Theory*

Arnab Kumar Mondal, Siba Smarak Panigrahi, Sai Rajeswar, K. Siddiqi, Siamak Ravanbakhsh

#157 *Reasoning with Latent Diffusion in Offline Reinforcement Learning*

Siddarth Venkatraman, Shivesh Khaitan, Ravi Tej Akella, John Dolan, Jeff Schneider, Glen Berseth

#160 *Provable and Practical: Efficient Exploration in Reinforcement Learning via Langevin Monte Carlo*

Haqee Ishfaq, Qingfeng Lan, Pan Xu, A. Rupam Mahmood, Doina Precup, Anima Anandkumar, Kamyar Azizzadenesheli

#164 *Consciousness-Inspired Spatio-Temporal Abstractions for Better Generalization in Reinforcement Learning*

Mingde Zhao, Safa Alver, Harm van Seijen, Romain Laroche, Doina Precup, Yoshua Bengio

#169 *Delta-AI: Local objectives for amortized inference in sparse graphical models*

Jean-Pierre R. Falet, Hae Beom Lee, Nikolay Malkin, Chen Sun, Dragos Secrieru, Dinghuai Zhang, Guillaume Lajoie, Yoshua Bengio

**Poster Session 7 -
10.45 a.m. (CET)**

#109 *Decoupling regularization from the action space*

Sobhan Mohammadpour, Emma Frejinger, Pierre-luc Bacon

#149 *Bridging State and History Representations: Understanding Self-Predictive RL*

Tianwei Ni, Benjamin Eysenbach, Erfan SeyedSalehi, Michel Ma, Clement Gehring, Aditya Mahajan, Pierre-luc Bacon

#155 *Intelligent Switching for Reset-Free RL*

Darshan Patil, Janarthanan Rajendran, Glen Berseth, Sarath Chandar

#203 *Less or More From Teacher: Exploiting Trilateral Geometry For Knowledge Distillation*

Chengming Hu, Haolun Wu, Xuan Li, Chen Ma, Xi Chen, Boyu Wang, Jun Yan, Xue Liu

#264 *Large Language Models as Generalizable Policies for Embodied Tasks*

Andrew Szot, Max Schwarzer, Harsh Agrawal, Bogdan Mazouze, Rin Metcalf, Katherine Metcalf, Walter Talbott, Natalie Mackraz, Devon Hjelm, Alexander Toshev

#265 *Motif: Intrinsic Motivation from Artificial Intelligence Feedback*

Martin Klissarov, P. D'Oro, Shagun Sodhani, Roberta Raileanu, Pierre-luc Bacon, Pascal Vincent, Amy Zhang, Mikael Henaff

**Orals -
3.45 p.m. (CET)**

Oral 8C

Mastering Memory Tasks with World Models

Mohammad Reza Samsami, Artem Zholus, Janarthanan Rajendran, Sarath Chandar

**Poster Session 8 -
4.30 p.m. (CET)**

#5 *Str2Str: A Score-based Framework for Zero-shot Protein Conformation Sampling*

Jiarui Lu, Bozitao Zhong, Zuobai Zhang, Jian Tang

#107 *The Curse of Diversity in Ensemble-Based Exploration*

Zhixuan Lin, P. D'Oro, Evgenii Nikishin, Aaron Courville

#149 *Expected flow networks in stochastic environments and two-player zero-sum games*

Marco Jiralerspong, Bilun Sun, Danilo Vucetic, Tianyu Zhang, Yoshua Bengio, Gauthier Gidel, Nikolay Malkin

#183 *Mastering Memory Tasks with World Models*

Mohammad Reza Samsami, Artem Zholus, Janarthanan Rajendran, Sarath Chandar

#191 *Piecewise Linear Parametrization of Policies: Towards Interpretable Deep Reinforcement Learning*

Maxime Wabartha, Joelle Pineau