

**Poster Session 1 -
11.30 a.m. (CET)**

#701 *Graph Positional and Structural Encoder*
Charles Guille-Escuret, Hiroki Naganuma, Kilian FATRAS, Ioannis Mitliagkas

#705 *CKGConv: General Graph Convolution with Continuous Kernels*

Liheng Ma, Soumyasundar Pal, Yitian Zhang, Jiaming Zhou, Yingxue Zhang, Mark Coates

#707 *EiG-Search: Generating Edge-Induced Subgraphs for GNN Explanation in Linear Time*

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

#1115 *Modeling Caption Diversity in Contrastive Vision-Language Pretraining*

Samuel Lavoie, Polina Kirichenko, Mark Ibrahim, Mahmoud Assran, Andrew Gordon Wilson, Aaron Courville, Nicolas Ballas

#1212 *Lookbehind-SAM: k steps back, 1 step forward*

Goncalo Mordido, Pranshu Malviya, Aristide Baratin, Sarath Chandar

#1403 *Motif: Intrinsic Motivation from Artificial Intelligence Feedback*

Learning to Play Atari in a World of Tokens

#2711 *Implicit meta-learning may lead language models to trust more reliable sources*

Dmitrii Krasheninnikov, Egor Krasheninnikov, Bruno Mlodozieniec, Tegan Maharaj, David Krueger

#2801 *A Persuasive Approach to Combating Misinformation*

Safwan Hossain, Andjela Mladenovic, Yiling Chen, Gauthier Gidel

#2905 *Position: Application-Driven Innovation in Machine Learning*

David Rolnick, Alan Aspuru-Guzik, Sara Beery, Bistra Dilkina, Priya L. Donti, Marzyeh Ghassemi, Hannah Kerner, Claire Monteleoni, Esther Rolf, Milind Tambe, Adam White

**Poster Session 2 -
1.30 p.m. (CET)**

#503 *Towards Modular LLMs by Building and Reusing a Library of LoRAs*

Oleksiy Ostapenko, Zhan Su, Edoardo Ponti, Laurent Charlin, Nicolas Le Roux, Matheus Pereira, Lucas Caccia, Alessandro Sordoni

#610 *WorkArena: How Capable are Web Agents at Solving Common Knowledge Work Tasks?*

Alexandre Drouin, Maxime Gasse, Massimo Caccia, Issam Hadj Laradji, Manuel Del Verme, Tom Marty, Léo Boisvert, Megh Thakkar, Quentin Cappart, David Vazquez, Nicolas Chapados, Alexandre Lacoste

#800 *SelfIE: Self-Interpretation of Large Language Model Embeddings*

Haozhe Chen, Carl Vondrick, Chengzhi Mao

#908 *Stochastic positional embeddings improve masked image modeling*

Amir Bar, Florian Bordes, Assaf Shocher, Mahmoud Assran, Pascal Vincent, Nicolas Ballas, Trevor Darrell, Amir Globerson, Yann LeCun

#1102 *Robust Data-driven Prescriptiveness Optimization*

Mehran Poursoltani, Érick Delage, Angelos Georgioui

#1112 *Sarah Frank-Wolfe: Methods for Constrained Optimization with Best Rates and Practical Features*

Aleksandr Beznosikov, David Dobre, Gauthier Gidel

#1117 *On PI Controllers for Updating Lagrange Multipliers in Constrained Optimization*

Motahareh Sohrabi, Juan Ramirez, Tianyue H. Zhang, Simon Lacoste-Julien, Jose Gallego-Posada

**Poster Session 2 -
1.30 p.m. (CET)**

#2313 *Position: On the Societal Impact of Open Foundation Models*

Sayash Kapoor, Rishi Bommasani, Kevin Klyman, Shayne Longpre, Ashwin Ramaswami, Peter Cihon, Aspen K Hopkins, Kevin Bankston, Stella Biderman, Miranda Bogen, Rumman Chowdhury, Alex Engler, Peter Henderson, Yacine Jernite, Seth Lazar, Stefano Maffulli, Alondra Nelson, Joelle Pineau, Aviya Skowron, Dawn Song, Victor Storchan, Daniel Zhang, Daniel E. Ho, Percy Liang, Arvind Narayanan

#2505 *Listenable Maps for Audio Classifiers*

Francesco Paissan, Mirco Ravanelli, Cem Subakan

#2704 *Position: Cracking the Code of Cascading Disparity Towards Marginalized Communities*

Golnoosh Farnadi, Mohammad Havaei, Negar Rostamzadeh

#2912 *Adaptive Accompaniment with ReaLchords*

Yusong Wu, Tim Cooijmans, Kyle Kastner, Adam Roberts, Ian Simon, Alexander Scarlatos, Chris Donahue, Cassie Tarakajian, Shayegan Omidshafiei, Aaron Courville, Pablo Samuel Castro, Natasha Jaques, Cheng-Zhi Anna Huang

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

Oral 2B - *Position: On the Societal Impact of Open Foundation Models*

Sayash Kapoor, Rishi Bommasani, Kevin Klyman, Shayne Longpre, Ashwin Ramaswami, Peter Cihon, Aspen K Hopkins, Kevin Bankston, Stella Biderman, Miranda Bogen, Rumman Chowdhury, Alex Engler, Peter Henderson, Yacine Jernite, Seth Lazar, Stefano Maffulli, Alondra Nelson, Joelle Pineau, Aviya Skowron, Dawn Song, Victor Storchan, Daniel Zhang, Daniel E. Ho, Percy Liang, Arvind Narayanan

Oral 2D - *Listenable Maps for Audio Classifiers*

Francesco Paissan, Mirco Ravanelli, Cem Subakan

Poster Session 3 -
11.30 a.m. (CET)
#311 Successor Features for Efficient Multi-Subject Controlled Text Generation

Meng Cao, Mehdi Fatemi, Jackie C. K. Cheung, Samira Shabania

#502 Don't be so negative! Score-based Generative Modeling with Oracle-assisted Guidance

Saeid Naderiparizi, Xiaoxuan Liang, Setareh Cohan, Berend Zwartsenberg, Frank Wood

#1411 Learning to Scale Logits for Temperature-Conditional GFlowNets

Minsu Kim, Joohwan Ko, Dinghuai Zhang, Ling Pan, Taeyoung Yun, Woo Chang Kim, Jinkyoo Park, Emmanuel Bengio, Yoshua Bengio

#2400 - Spotlight Beyond the Norms: Detecting Prediction Errors in Regression Models

Andres Altieri, Marco Romanelli, Georg Pichler, Florence Alberge, Pablo Piantanida

#2515 - Spotlight Faithfulness Measurable Masked Language Models

Andreas Madsen, Siva Reddy, Sarath Chandar, Milind Tambe, Adam White

Poster Session 4 -
1.30 p.m. (CET)
#407 - Memory Efficient Neural Processes via Constant Memory Attention Block

Leo Feng, Frederick Tung, Hossein Hajimir-sadeghi, Yoshua Bengio, Mohamed Osama Ahmed

#417 - Nearest Neighbour Score Estimators for Diffusion Generative Models

Matthew Niedoba, Dylan Green, Saeid Naderiparizi, Vasileios Lioutas, Jonathan Wilder Lavington, Xiaoxuan Liang, Yunpeng Liu, Ke Zhang, Setareh Dabiri, Adam Ścibior, Berend Zwartsenberg, Frank Wood

Poster Session 4 -
1.30 p.m. (CET)
#816 Harmony in Diversity: Merging Neural Networks with Canonical Correlation Analysis

Stefan Horoi, Albert Manuel Orozco Camacho, Eugene Belilovsky, Guy Wolf

#1005 Experts Don't Cheat: Learning What You Don't Know By Predicting Pairs

Daniel D. Johnson, Danny Tarlow, David Duvenaud, Chris J. Maddison

#1010 No Wrong Turns: The Simple Geometry Of Neural Networks Optimization Paths

Charles Guille-Escuret, Hiroki Naganuma, Kilian FATRAS, Ioannis Mitliagkas

#1200 PcLast: Discovering Plannable Continuous Latent States

Anurag Koul, Shivakanth Sujit, Shaoru Chen, Ben Evans, Lili Wu, Byron Xu, Rajan Chari, Riashat Islam, Raihan Seraj, Yonathan Efroni, Lekan Molu, Miro Dudik, John Langford, Alex Lamb

#1308 In value-based deep reinforcement learning, a pruned network is a good network

Johan Samir Obando Ceron, Aaron Courville, Pablo Samuel Castro

#1311 Stop Regressing: Training Value Functions via Classification for Scalable Deep RL

Jesse Farebrother, Jordi Orbay, Quan Vuong, Adrien Ali Taiga, Yevgen Chebotar, Ted Xiao, Alex Irpan, Sergey Levine, Pablo Samuel Castro, Aleksandra Faust, Aviral Kumar, Rishabh Agarwal

#1405 - Spotlight A Distributional Analogue to the Successor Representation

Harley Wiltzer, Jesse Farebrother, Arthur Gretton, Yunhao Tang, Andre Barreto, Will Dabney, Marc G. Bellemare, Mark Rowland

#1406 - Spotlight Estimating Unknown Population Sizes Using the Hypergeometric Distribution

Liam Hodgson, Danilo Bzdok

Poster Session 2 -
1.30 p.m. (CET)
#1408 A Computational Framework for Solving Wasserstein Lagrangian Flows

Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, qiang liu, Alireza Makhzani

#1414 Improving Gradient-Guided Nested Sampling for Posterior Inference

Golnoosh Farnadi, Mohammad Havaei, Negar Rostamzadeh

#1415 Iterated Denoising Energy Matching for Sampling from Boltzmann Densities

Tara Akhound-Sadegh, Jarrid Rector-Brooks, Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, Alexander Tong

#1806 Randomized Confidence Bounds for Stochastic Partial Monitoring

Maxime Heuillet, Ola Ahmad, Audrey Durand

#2308 Stealing part of a production language model

Nicholas Carlini, Daniel Paleka, Krishnamurthy Dj Dvijotham, Thomas Steinke, Jonathan Hayase, A. Feder Cooper, Katherine Lee, Matthew Jagielski, Milad Nasr, Arthur Conmy, Eric Wallace, David Rolnick, Florian Tramèr

Orals -
4.30 p.m. (CET)
Oral 4A - Stop Regressing: Training Value Functions via Classification for Scalable Deep RL

Jesse Farebrother, Jordi Orbay, Quan Vuong, Adrien Ali Taiga, Yevgen Chebotar, Ted Xiao, Alex Irpan, Sergey Levine, Pablo Samuel Castro, Aleksandra Faust, Aviral Kumar, Rishabh Agarwal

Oral 4C - Stealing part of a production language model

Nicholas Carlini, Daniel Paleka, Krishnamurthy Dj Dvijotham, Thomas Steinke, Jonathan Hayase, A. Feder Cooper, Katherine Lee, Matthew Jagielski, Milad Nasr, Arthur Conmy, Eric Wallace, David Rolnick, Florian Tramèr

Orals - 10.30 a.m. (CET)

Oral 5B - *Information Complexity of Stochastic Convex Optimization: Applications to Generalization and Memorization*

Idan Attias, Gintare Karolina Dziugaite, MAHDI HAGHIFAM, Roi Livni, Daniel M. Roy

Oral 5 - *High-Probability Convergence for Composite and Distributed Stochastic Minimization and Variational Inequalities with Heavy-Tailed Noise*

Eduard Gorbunov, Abdurakhmon Sadiev, Marina Danilova, Samuel Horváth, Gauthier Gidel, Pavel Dvurechensky, Alexander Gasnikov, Peter Richtárik

Oral 5E - *Discovering environments with XRM*

Mohammad Pezeshki, Diane Bouchacourt, Mark Ibrahim, Nicolas Ballas, Pascal Vincent, David Lopez-Paz

Poster Session 5 - 11.30 a.m. (CET)

#403 - *Discovering environments with XRM*

Mohammad Pezeshki, Diane Bouchacourt, Mark Ibrahim, Nicolas Ballas, Pascal Vincent, David Lopez-Paz

#1014 - *High-Probability Convergence for Composite and Distributed Stochastic Minimization and Variational Inequalities with Heavy-Tailed Noise*

Eduard Gorbunov, Abdurakhmon Sadiev, Marina Danilova, Samuel Horváth, Gauthier Gidel, Pavel Dvurechensky, Alexander Gasnikov, Peter Richtárik

#1201 - *Learning to Reach Goals via Diffusion*

Vineet Jain, Siamak Ravanbakhsh

#1203 - *Do Transformer World Models Give Better Policy Gradients?*

Michel Ma, Tianwei Ni, Clement Gehring, Pierluca D'Oro, Pierre-luc Bacon

Poster Session 5 - 11.30 a.m. (CET)

#1902 - *Layerwise Proximal Replay: A Proximal Point Method for Online Continual Learning*

Jinsoo Yoo, Yunpeng Liu, Frank Wood, Geoff Pleiss

#2400 - *Information Complexity of Stochastic Convex Optimization: Applications to Generalization and Memorization*

Idan Attias, Gintare Karolina Dziugaite, MAHDI HAGHIFAM, Roi Livni, Daniel M. Roy

Poster Session 6 - 1.30 p.m. (CET)

#305 *Autoformalizing Euclidean Geometry*

Logan Murphy, Kaiyu Yang, Jialiang Sun, Zhaoyu Li, Anima Anandkumar, Xujie Si

#317 - Spotlight *WebLINX: Real-World Website Navigation with Multi-Turn Dialogue*

Xing Han Lù, Zdeněk Kasner, Siva Reddy

#400 *Unsupervised Concept Discovery Mitigates Spurious Correlations*

Md Rifat Arefin, Yan Zhang, Aristide Baratin, Francesco Locatello, Irina Rish, Dianbo Liu, Kenji Kawaguchi

#708 - Spotlight *Nash Learning from Human Feedback*

Remi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Zhaohan Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Côme Fiegel, Andrea Michi, Marco Selvi, Sertan Girgin, Nikola Momchev, Olivier Bachem, Daniel J Mankowitz, Doina Precup, Bilal Piot

#910 - Spotlight *Tensor Decomposition Perspective on Second-order RNNs*

Maude Lizaire, Michael Rizvi-Martel, Marwan Gamal, Guillaume Rabusseau

Poster Session 6 - 1.30 p.m. (CET)

#911 *Interacting Diffusion Processes for Event Sequence Forecasting*

Mai Zeng, Florence Regol, Mark Coates

#1115 - Spotlight *Code as Reward: Empowering Reinforcement Learning with VLMs*

David Venuto, Mohammad Sami Nur Islam, Martin Klissarov, Doina Precup, Sherry Yang, Ankit Anand

#1207 - Spotlight *Mixtures of Experts Unlock Parameter Scaling for Deep RL*

Johan Samir Obando Ceron, Ghada Sokar, Timon Willi, Clare Lyle, Jesse Farebrother, Jakob Nicolaus Foerster, Gintare Karolina Dziugaite, Doina Precup, Pablo Samuel Castro

#1309 - *All-in-one simulation-based inference*

Manuel Gloeckler, Michael Deistler, Christian Dietrich Weilbach, Frank Wood, Jakob H. Macke

Orals - 4.30 p.m. (CET)

Oral 6F - *All-in-one simulation-based inference*

Manuel Gloeckler, Michael Deistler, Christian Dietrich Weilbach, Frank Wood, Jakob H. Macke