Mila @ NeurIPS 2025 / 12.03.2025

Follow Mila researchers at NeurIPS 2025
Suivez les chercheur·se·s de Mila à NeurIPS 2025

Mila
Institut québécois d'intelligence artificielle
Quebec Artificial Intelligence Institute

Mila's Booth / kiosque de Mila

11:00 AM Pooneh Mousavi - *Learning Audio Representations for Multimodal Models*

1:00 PM Martin Weiss - *Multi-agent systems and tools for ML researchers*

Poster Session 1 - 11 AM - 2 PM

SAN DIEGO

#2112 *Caption This, Reason That: VLMs Caught in the Middle* - Zihan Weng, Lucas Gomez, Taylor Whittington Webb, Pouya Bashivan

#706 *New Perspectives on the Polyak Stepsize: Surrogate Functions and Negative Results* - Francesco Orabona, Ryan D'Orazio

#1400 *From Dormant to Deleted: Tamper-Resistant Unlearning Through Weight-Space Regularization* - Shoaib Ahmed Siddiqui, Adrian Weller, David Krueger, Gintare Karolina Dziugaite, Michael Curtis Mozer, Eleni Triantafillou

#1807 *FocalCodec: Low-Bitrate Speech Coding via Focal Modulation Networks* - Luca Della Libera, Francesco Paissan, Yusuf Cem Sübakan, Mirco Ravanelli

#2000 *Generalizable, real-time neural decoding with hybrid state-space models* - Avery Hee-Woon Ryoo, Nanda H Krishna, Ximeng Mao, Mehdi Azabou, Eva L Dyer, Matthew G Perich, Guillaume Lajoie

#3407 *Reducing the Probability of Undesirable Outputs in Language Models Using Probabilistic Inference* - Stephen Zhao, Aidan Li, Rob Brekelmans, Roger Grosse

#3514 *Adaptive Inference-Time Scaling via Cyclic Diffusion Search* - Gyubin Lee, Bao N Nguyen Truong, Jaesik Yoon, Dongwoo Lee, Minsu Kim, Yoshua Bengio, Sungjin Ahn

#3803 *AlignVLM: Bridging Vision and Language Latent Spaces for Multimodal Document Understanding* - Ahmed Masry, Juan A. Rodriguez, Tianyu Zhang, Suyuchen Wang, Chao Wang, Aarash Feizi, Akshay Kalkunte Suresh, Abhay Puri, Xiangru Jian, Pierre-Andre Noel, Sathwik Tejaswi Madhusudhan, Marco Pedersoli, Bang Liu, Nicolas Chapados, Yoshua Bengio, Enamul Hoque, Christopher Pal, Issam Hadj Laradji, David Vazquez, Perouz Taslakian, Spandana Gella, Sai Rajeswar

#3917 *Geometry-Aware Edge Pooling for Graph Neural Networks* - Katharina Limbeck, Lydia Mezrag, Guy Wolf, Bastian Rieck

#4205 *Transforming Generic Coder LLMs to Effective Binary Code Embedding Models for Similarity Detection* - Katharina Limbeck, Lydia Mezrag, Guy Wolf, Bastian Rieck

#4300 *Increasing the Utility of Synthetic Images through Chamfer Guidance* - Nicola Dall'Asen, Xiaofeng Zhang, Reyhane Askari Hemmat, Melissa Hall, Jakob Verbeek, Adriana Romero, Michal Drozdzal

#5414 *Mixture-of-Recursions: Learning Dynamic Recursive Depths for Adaptive Token-Level Computation* - Sangmin Bae, Yujin Kim, Reza Bayat, Sungnyun Kim, Jiyoun Ha, Tal Schuster, Adam Fisch, Hrayr Harutyunyan, Ziwei Ji, Aaron C. Courville, Se-Young Yun

DATASET AND BENCHMARK TRACK

#1513 *Open-Insect: Benchmarking Open-Set Recognition of Novel Species in Biodiversity Monitoring* - Yuyan Chen, Nico Lang, B. Christian Schmidt, Aditya Jain, Yves Basset, Sara Beery, Maxim Larrivée, David Rolnick

#403 *Meta-World+: An Improved, Standardized, RL Benchmark* - Reginald McLean, Evangelos Chatzaroulas, Luc McCutcheon, Frank Röder, Tianhe Yu, Zhanpeng He, K.R. Zentner, Ryan Julian, J K Terry, Isaac Woungang, Nariman Farsad, Pablo Samuel Castro

MEXICO

The Promise of RL for Autoregressive Image Editing - Saba Ahmadi, Rabiul Awal, Ankur Sikarwar, Amirhossein Kazemnejad, Ge Ya Luo, Juan A. Rodriguez, Sai Rajeswar, Siva Reddy, Christopher Pal, Benno Krojer, Aishwarya Agrawal

TRUST: Test-Time Refinement using Uncertainty-Guided SSM Traverses - Sahar Dastani, Ali Bahri, Gustavo Adolfo Vargas Hakim, Moslem Yazdanpanah, Mehrdad Noori, David Osowiechi, Samuel Barbeau, Ismail Ben Ayed, Herve Lombaert, Christian Desrosiers

PointMAC: Meta-Learned Adaptation for Robust Test-Time Point Cloud Completion - Linlian Jiang, Rui Ma, Li Gu, Ziqiang Wang, Xinxin Zuo and Yang Wang

Poster Session 2

SAN DIEGO - 4:30PM to 7:30PM

#1016 *Beyond Scalar Rewards: An Axiomatic Framework for Lexicographic MDPs* - Mehran Shakerinava, Siamak Ravanbakhsh, Adam Oberman

#3201 *Dimension-adapted Momentum Outscales SGD* - Damien Ferbach, Katie Everett, Gauthier Gidel, Elliot Paquette, Courtney Paquette

#908 *Understanding Adam Requires Better Rotation Dependent Assumptions* - Tianyue H. Zhang, Lucas Maes, Alan Milligan, Alexia Jolicoeur-Martineau, Ioannis Mitliagkas, Damien Scieur, Simon Lacoste-Julien, Charles Guille-Escuret

#909 *ACCO: Accumulate While You Communicate for Communication-Overlapped Sharded LLM Training* - Adel Nabli, Louis Fournier, Pierre ERBACHER, Louis Serrano, Eugene Belilovsky, Edouard Oyallon

#4100 *Rendering-Aware Reinforcement Learning for Vector Graphics Generation - Juan A. Rodriguez, Haotian Zhang, Abhay Puri, Rishav Pramanik, Aarash Feizi, Pascal Wichmann, Arnab Kumar Mondal, Mohammad Reza Samsami, Rabiul Awal, Perouz Taslakian, Spandana Gella, Sai Rajeswar, David Vazquez, Christopher Pal, Marco Pedersoli*

#4919 *Bringing SAM to new heights: Leveraging elevation data for tree crown segmentation from drone imagery - Mélisande Teng, Arthur Ouaknine, Etienne Laliberté, Yoshua Bengio, David Rolnick, Hugo Larochelle*

DATASET AND BENCHMARK TRACK

#3802 *MiNT: Multi-Network Transfer Benchmark for Temporal Graph Learning* - Kiarash Shamsi, Tran Gia Bao Ngo, Razieh Shirzadkhani, Shenyang Huang, Farimah Poursafaei, Poupak Azad, Reihaneh Rabbany, Baris Coskunuzer, Guillaume Rabusseau, Cuneyt Gurcan Akcora

POSITION PAPER TRACK

#1313 *Neither Valid Nor Reliable? Investigating the Use of LLMs as Judges* - Khaoula Chehbouni, Mohammed Haddou, Jackie CK Cheung, Golnoosh Farnadi

MEXICO - 2:30PM to 7:30PM

PCA++: How Uniformity Induces Robustness to Background Noise in Contrastive Learning - Mingqi Wu · Qiang Sun · Archer Yang

Tight Lower Bounds and Improved Convergence in Performative Prediction - Pedram Khorsandi, Rushil Gupta, Mehrnaz Mofakhami, Simon Lacoste-Julien, Gauthier Gidel

Mila @ NeurIPS 2025 / 12.04.2025

Follow Mila researchers at NeurIPS 2025
Suivez les chercheur·se·s de Mila à NeurIPS 2025

Mila
Institut québécois d'intelligence artificielle
Quebec Artificial Intelligence Institute

Mila's Booth / kiosque de Mila

11:00 AM Juan Ramirez - *Feasible Learning*

1:00 PM Padideh Nouri - *Relative Trajectory Balance is equivalent to Trust-PCL*

Poster Session 3 - 11 AM - 2 PM

SAN DIEGO

- **#2407** *A Smooth Sea Never Made a Skilled SAILOR: Robust Imitation via Learning to Search* - Arnav Kumar Jain, Vibhakar Mohta, Subin Kim, Atiksh Bhardwaj, Juntao Ren, Yunhai Feng, Sanjiban Choudhury, Gokul Swamy
- **#3109** *Is the acquisition worth the cost? Surrogate losses for Consistent Two-stage Classifiers* - Florence Regol, Joseph Cotnareanu, Theodore Glavas, Mark J. Coates

#503 *Epistemic Uncertainty Estimation in Regression Ensemble Models with Pairwise Epistemic Estimators* - Lucas Berry, David Meger

#512 *State Entropy Regularization for Robust Reinforcement Learning* - Yonatan Ashlag, Uri Koren, Mirco Mutti, Esther Derman, Pierre-Luc Bacon, Shie Mannor

#1112 *Detecting High-Stakes Interactions with Activation Probes* - Alex McKenzie, Urja Pawar, Phil Blandfort, William Bankes, David Krueger, Ekdeep Singh Lubana, Dmitrii Krasheninnikov

#1201 *Uncovering a Universal Abstract Algorithm for Modular Addition in Neural Networks* - Gavin McCracken, Gabriela Moises-cu-Pareja, Vincent Létourneau, Doina Precup, Jonathan Love

#1614 *Amortized Sampling with Transferable Normalizing Flows* - Charlie B. Tan, Majdi Hassan, Leon Klein, Saifuddin Syed, Dominique Beaini, Michael M. Bronstein, Alexander Tong, Kirill Neklyudov

#1705 *Causal Climate Emulation with Bayesian Filtering* - Sebastian Hickman, Ilija Trajković, Julia Kaltenborn, Francis Pelletier, Alexander T Archibald, Yaniv Gurwicz, Peer Nowack, David Rolnick, Julien Boussard

#3917 *Geometry-Aware Edge Pooling for Graph Neural Networks* - Katharina Limbeck, Lydia Mezrag, Guy Wolf, Bastian Rieck

#1708 *Scalable and Cost-Efficient de Novo Template-Base Molecular Generation* - Piotr Gaiński, Oussama Boussif, Andrei Rekesh, Dmytro Shevchuk, Ali Parviz, Mike Tyers, Robert A. Batey, Michał Koziarski

#2201 *Dimensionality Mismatch Between Brains and Artificial Neural Networks* - Santiago Galella, Maren Wehrheim, Matthias Kaschube

#2715 *Random Forest Autoencoders for Guided Representation Learning* - Adrien Aumon, Shuang Ni, Myriam Lizotte, Guy Wolf, Kevin R. Moon, Jake S. Rhodes

#5209 *Adaptive Quantization in Generative Flow Networks for Probabilistic Sequential Prediction* - Nadhir Hassen, Zhen Zhang, Johan W. Verjans

#5409 *C3PO: Optimized Large Language Model Cascades with Probabilistic Cost Constraints for Reasoning* - Antonios Valkanas, Soumyasundar Pal, Pavel Rumiantsev, Yingxue Zhang, Mark J. Coates

POSITION PAPER TRACK

#1008 *Rigor in AI: Doing Rigorous AI Work Requires a Broader, Responsible AI-Informed Conception of Rigor* - Alexandra Olteanu, Su Lin Blodgett, Agathe Balayn, Angelina Wang, Fernando Díaz, Flavio Calmon, Margaret Mitchell, Michael Ekstrand, Reuben Binns, Solon Barocas

#403 *Meta-World+: An Improved, Standardized, RL Benchmark* - Reginald McLean, Evangelos Chatzaroulas, Luc McCutcheon, Frank Röder, Tianhe Yu, Zhanpeng He, K.R. Zentner, Ryan Julian, J K Terry, Isaac Woungang, Nariman Farsad, Pablo Samuel Castro

MEXICO

Learning to Solve Complex Problems via Dataset Decomposition - Wanru Zhao, Lucas Caccia, Zhengyan Shi, Minseon Kim, Weijia Xu, Xingdi Yuan, Alessandro Sordoni, Marc-Alexandre Côté

Poster Session 4

SAN DIEGO - 4:30PM to 7:30PM

- **#512** *Plasticity as the Mirror of Empowerment* - David Abel, Michael Bowling, Andre Barreto, Will Dabney, Shi Dong, Steven Stenberg Hansen, Anna Harutyunyan, Khimya Khetarpal, Clare Lyle, Razvan Pascanu, Georgios Piliouras, Doina Precup, Jonathan Richens, Mark Rowland, Tom Schaul, Satinder Singh

#314 *Deep RL Needs Deep Behavior Analysis: Exploring Implicit Planning by Model-Free Agents in Open-Ended Environments* - Riley Simmons-Edler, Ryan P Badman, Felix Baastad Berg, Raymond Chua, John J Vastola, Joshua Lunger, William Qian, Kanaka Rajan

#407 *How to Train Your LLM Web Agent: A Statistical Diagnosis* - Dheeraj Vattikonda, Santhoshi Ravichandran, Emiliano Penaloza, Hadi Nekoei, Thibault Le Sellier de Chezelles, Megh Thakkar, Nicolas Gontier, Miguel Muñoz-Mármol, Sahar Omidi Shayegan, Stefania Raimondo, Xue Liu, Alexandre Drouin, Alexandre Piché, Laurent Charlin, Alexandre Lacoste, Massimo Caccia

#1710 *RETRO SYNFLOW: Discrete Flow Matching for Accurate and Diverse Single-Step Retrosynthesis* - Robin Yadav, Qi Yan, Guy Wolf, Avishek Joey Bose, Renjie Liao

#1900 *System-1.5 Reasoning: Traversal in Language and Latent Spaces with Dynamic Shortcuts - Xiaoqiang Wang, Suyuchen Wang, Yun Zhu, Bang Liu*

#2209 *Improving Energy Natural Gradient Descent through Woodbury, Momentum, and Randomization - Andrés Guzmán-Cordero, Felix Dangel, Gil Goldshlager, Marius Zeinhofer*

#2708 *Learning Task-Agnostic Representations through Multi-Teacher Distillation - Philippe Formont, Maxime DARRIN, Banafsheh Karimian, Eric Granger, Jackie CK Cheung, Ismail Ben Ayed, Mohammadhadi Shateri, Pablo Piantanida*

DATASET AND BENCHMARK TRACK

#1700 *GreenHyperSpectra: A multi-source hyperspectral dataset for global vegetation trait prediction* - Eya Cherif, Arthur Ouaknine, Luke A. Brown, Phuong D. Dao, Kyle R Kovach, Bing Lu, Daniel Mederer, Hannes Feilhauer, Teja Kattenborn, David Rolnick

Mila @ NeurIPS 2025 / 12.05.2025

Follow Mila researchers at NeurIPS 2025
Suivez les chercheur·se·s de Mila à NeurIPS 2025

Mila
Institut québécois d'intelligence artificielle
Quebec Artificial Intelligence Institute

Poster Session 5 - 11 AM - 2 PM

SAN DIEGO

- #310 *Stable Gradients for Stable Learning at Scale in Deep Reinforcement Learning* - Roger Creus Castanyer, Johan Samir Obando Ceron, Li Li, Pierre-Luc Bacon, Glen Berseth, Aaron C. Courville, Pablo Samuel Castro
- #3609 *Fast Monte Carlo Tree Diffusion: 100× Speedup via Parallel and Sparse Planning* - Jaesik Yoon, Hyeonseo Cho, Yoshua Bengio, Sungjin Ahn
- #3901 *Distributional Training Data Attribution: What do Influence Functions Sample?* - Bruno Mlodozeniec, Isaac Reid, Samuel Power, David Krueger, Murat A Erdogdu, Richard E. Turner, Roger Baker Grosse

#206 *Tapered Off-Policy REINFORCE - Stable and efficient reinforcement learning for large language models* - Nicolas Roux, Bellemare Marc-Emmanuel, Jonathan Lebensold, Arnaud Bergeron, Joshua Greaves, Alexandre Fréchette, Carolyne Pelletier, Eric Thibodeau-Laufer, Sándor Tóth, Sam Work

#214 *Measure gradients, not activations! Enhancing neuronal activity in deep reinforcement learning* - Jiashun Liu, Zihao Wu, Johan Samir Obando Ceron, Pablo Samuel Castro, Aaron C. Courville, Ling Pan

#1106 *From Noise to Narrative: Tracing the Origins of Hallucinations in Transformers* - Praneet Suresh, Jack Stanley, Sonia Joseph, Luca Scimeca, Danilo Bzdok

#2014 *REVE: A Foundation Model for EEG - Adapting to Any Setup with Large-Scale Pretraining on 25,000 Subjects* - Yassine El Ouahidi, Jonathan Lys, Philipp Thölke, Nicolas Farrugia, Bastien Pasdeloup, Vincent Gripon, Karim Jerbi, Giulia Lioi

#2107 *Know Thyself by Knowing Others: Learning Neuron Identity from Population Context* - Vinam Arora, Divyansha Lachi, Ian Jarratt Knight, Mehdi Azabou, Blake Aaron Richards, Cole Lincoln Hurwitz, Josh Siegle, Eva L Dyer

#3502 *Diffusion Tree Sampling: Scalable inference-time alignment of diffusion models* - Vineet Jain, Kusha Sareen, Mohammad Pedramfar, Siamak Ravanbakhsh

#3505 *Compositional Discrete Latent Code for High Fidelity, Productive Diffusion Models* - Samuel Lavoie, Michael Noukhovitch, Aaron C. Courville

#3414 *Overcoming Long-Context Limitations of State-Space Models via Context-Dependent Sparse Attention* - Zhihao Zhan, Jianan Zhao, Zhaocheng Zhu, Jian Tang

#4010 *Discovering Data Structures: Nearest Neighbor Search and Beyond* - Omar Salemohamed, Laurent Charlin, Shivam Garg, Vatsal Sharan, Gregory Valiant

#4314 *Entropy Rectifying Guidance for Diffusion and Flow Models* - Tariq Berrada, Adriana Romero, Michal Drozdzal, Jakob Verbeek, Karteek Alahari

#4315 *Discovering Latent Graphs with GFlowNets for Diverse Conditional Image Generation* - Bailey Trang, Parham Saremi, Alan Q. Wang, Fangrui Huang, Zahra Tehrani Nasab, Amar Kumar, Tal Arbel, Li Fei-Fei, Ehsan Adeli

DATASET AND BENCHMARK TRACK

#311 *Generating Creative Chess Puzzles* - Xidong Feng, Vivek Veeriah, Marcus Chiam, Michael D Dennis, Federico Barbero, Johan Obando-Ceron, Jiaxin Shi, Satinder Singh, Shaobo Hou, Nenad Tomasev, Tom Zahavy

#4913 *OpenLex3D: A Tiered Benchmark for Open-Vocabulary 3D Scene Representations* - Christina Kassab, Sacha Morin, Martin Büchner, Matias Mattamala, Kumaraditya Gupta, Abhinav Valada, Liam Paull, Maurice Fallon

MEXICO

Majority of the Bests: Improving Best-of-N via Bootstrapping - Amin Rakhsha, Kanika Madan, Tianyu Zhang, Amir-massoud Farahmand, Amir Khasahmadi

Poster Session 6

SAN DIEGO - 4:30PM to 7:30PM

- #1604 *Progressive Inference-Time Annealing of Diffusion Models for Sampling from Boltzmann Densities* - Tara Akhound-Sadegh, Jungyoon Lee, Avishek Joey Bose, Valentin De Bortoli, Arnaud Doucet, Michael M. Bronstein, Dominique Beaini, Siamak Ravanbakhsh, Kirill Neklyudov, Alexander Tong
- #2603 *Causal Differentiating Concepts: Interpreting LM Behavior via Causal Representation Learning* - Navita Goyal, Hal Daumé III, Alexandre Drouin, Dhanya Sridhar

#304 *Trajectory Balance with Asynchrony: Decoupling Exploration and Learning for Fast, Scalable LLM Post-Training* - Brian R. Bartoldson, Siddarth Venkatraman, James Diffenderfer, Moksh Jain, Tal Ben-Nun, Seanie Lee, Minsu Kim, Johan Samir Obando Ceron, Yoshua Bengio, Bhavya Kailkhura

#2006 *POCO: Scalable Neural Forecasting through Population Conditioning* - Yu Duan, Hamza Tahir Chaudhry, Misha B. Ahrens, Christopher D Harvey, Matthew G Perich, Karl Deisseroth, Kanaka Rajan

#2103 *Energy Loss Functions for Physical Systems* - Sékou-Oumar Kaba, Kusha Sareen, Daniel Levy, Siamak Ravanbakhsh

#2502 *Tracing the representation geometry of language models from pretraining to post-training* - Melody Zixuan Li, Kumar Krishna Agrawal, Arna Ghosh, Komal Kumar Teru, Adam Santoro, Guillaume Lajoie, Blake Aaron Richards

#3102 *Planning and Learning in Average Risk-aware MDPs* - Weikai Wang, Erick Delage

#3107 *Convergence Theorems for Entropy-Regularized and Distributional Reinforcement Learning* - Yash Jhaveri, Harley Wiltzer, Patrick Shafto, Marc G. Bellemare, David Meger

#4606 *Object-centric Binding in Contrastive Language-Image Pretraining* - Rim Assouel, Pietro Astolfi, Florian Bordes, Michal Drozdzal, Adriana Romero

DATASET AND BENCHMARK TRACK

- #4413 *THUNDER: Tile-level Histopathology image UNDERstanding benchmark* - Pierre Marza, Leo Fillioux, Sofiène Boutaj, KUNAL MAHATHA, Christian Desrosiers, Pablo Piantanida, Jose Dolz, Stergios Christodoulidis, Maria Vakalopoulou

#213 *NAVIX: Scaling MiniGrid Environments with JAX* - Eduardo Pignatelli, Jarek Luca Liesen, Robert Tjarko Lange, Chris Lu, Pablo Samuel Castro, Laura Toni