Mila @ ICLR 2026
23.04.2026

Follow Mila researchers at ICLR 2026
Suivez les chercheuses et chercheurs de Mila à ICLR 2026

Mila
Institut québécois d'intelligence artificielle
Quebec Artificial Intelligence Institute

Poster Session 1 - 10.30 a.m. (BRT)

P3-#218 *DisTaC: Conditioning Task Vectors via Distillation for Robust Model Merging*
Kotaro Yoshida, Yuji Naraki, Takafumi Horie, Ryotaro Shimizu, Ioannis Mitliagkas, Hiroki Naganuma

P3-#803 *GraphOmni: A Comprehensive and Extensible Benchmark Framework for Large Language Models on Graph-theoretic Tasks*
Hao Xu, Xiangru Jian, Xinjian Zhao, Wei Pang, Chao Zhang, Suyuchen Wang, Qixin ZHANG, Zhengyuan Dong, Joao Monteiro, Bang Liu, Qiuzhuang Sun, Tianshu Yu

P3-#907 *Robust Fine-Tuning from Non-Robust Pretrained Models: Mitigating Suboptimal Transfer With Epsilon-Scheduling*
Jonas Ngnawe, Maxime Heuillet, Sabyasachi Sahoo, Yann Pequignot, Ola Ahmad, Audrey Durand, Frederic Precioso, Christian Gagné

P3-#1005 *Towards All-Atom Foundation Models for Biomolecular Binding Affinity Prediction*
Liang Shi, Zuobai Zhang, Huiyu Cai, Santiago Miret, Zhi Yang, Jian Tang

P3-#1101 *Efficient Regression-based Training of Normalizing Flows for Boltzmann Generators*
Danyal Rehman, Oscar Davis, Jiarui Lu, Jian Tang, Michael Bronstein, Yoshua Bengio, Alexander Tong, Joey Bose

P3-#1721 *Unraveling the Complexity of Memory in RL Agents: an Approach for Classification and Evaluation*
Egor Cherepanov, Nikita Kachaev, Artem Zholus, Alexey Kovalev, Aleksandr Panov

P4-#3906 *Model Collapse Is Not a Bug but a Feature in Machine Unlearning for LLMs*
Yan Scholten, Sophie Xhonneux, Leo Schwinn, Stephan Günnemann

P4-#3918 *Landscape of Thoughts: Visualizing the Reasoning Process of Large Language Models*
(Andrew) Zhanke Zhou, Zhaocheng Zhu, Xuan Li, Mikhail Galkin, Xiao Feng, Sanmi Koyejo, Jian Tang, Bo Han

P4-#4406 *A Derandomization Framework for Structure Discovery: Applications in Neural Networks and Beyond*
Nikos Tsikouras, Yorgos Pantis, Ioannis Mitliagkas, Christos Tzamos

P4-#4709 *SHAPO: Sharpness-Aware Policy Optimization for Safe Exploration*
Kaustubh Mani, Yann Pequignot, Vincent Mai, Liam Paull

P4-#4805 *ARM-FM: Automated Reward Machines via Foundation Models for Compositional Reinforcement Learning*
Roger Creus Castanyer, Faisal Mohamed, Pablo Samuel Castro, Cyrus Neary, Glen Berseth

Oral Session 1 - 10.30 a.m. (BRT)

Oral 1E On The Surprising Effectiveness of a Single Global Merging in Decentralized Learning
Tongtian Zhu, Tianyu Zhang, Mingze Wang, Zhanpeng Zhou, Can Wang

Poster Session 2 - 3.15 p.m. (BRT)

P3-#914 On The Surprising Effectiveness of a Single Global Merging in Decentralized Learning
Tongtian Zhu, Tianyu Zhang, Mingze Wang, Zhanpeng Zhou, Can Wang

P3-#609 *Setting the Record Straight on Transformer Oversmoothing*
Gbètondji J-S Dovonon, Michael Bronstein, Matt J. Kusner

P3-#1204 *Contractive Diffusion Policies: Robust Action Diffusion via Contractive Score-Based Sampling with Differential Equations*
Amin Soleimani Abyaneh, Charlotte Morissette, Mohamad H. Danesh, Anas Houssaini, David Meger, Gregory Dudek, Hsiu-Chin Lin

P3-#1705 *Grounding Computer Use Agents on Human Demonstrations*
Aarash Feizi, Shravan Nayak, Xiangru Jian, Kevin Qinghong Lin, Kaixin Li, Rabiul Awal, Xing Han Lu, Johan S Obando Ceron, Juan A. Rodriguez, Nicolas Chapados, David Vazquez, Adriana Romero-Soriano, Reihaneh Rabbany, Perouz Taslakian, Christopher Pal, Spandana Gella, Sai Rajeswar Mudumba

P3-#1815 *Hierarchical Value-Decomposed Offline Reinforcement Learning for Whole-Body Control*
Zhilong Zhang, Yunpeng Mei, Xinghao Du, Hongjie Cao, Haonan Wang, Pengyuan Min, Chenyu Wang, Pengfei Chen, Chenbo Xin, Yijie Wang, Wenyu Luo, Yihao Sun, Yidi Wang, Lei Yuan, Gang Wang, Yang Yu

P3-#2008 *Asymmetric Proximal Policy Optimization: mini-critics boost LLM reasoning*
Jiashun Liu, Johan S Obando Ceron, Han Lu, Yancheng He, Weixun Wang, wenbo su, Bo Zheng, Pablo Samuel Castro, Aaron Courville, Ling Pan

P4-#3114 *The Intricate Dance of Prompt Complexity, Quality, Diversity and Consistency in T2I Models*
Zhang Xiaofeng, Aaron Courville, Michal Drozdzal, Adriana Romero-Soriano

P4-#3701 *SelvaBox: A high-resolution dataset for tropical tree crown detection*
Hugo Baudchon, Arthur Ouaknine, Martin Weiss, Mélisande Teng, Thomas Walla, Antoine Caron-Guay, Christopher Pal, Etienne Laliberté

P4-#3805 *Unsupervised Representation Learning for 3D Mesh Parameterization with Semantic and Visibility Objectives*
Amirhossein Zamani, Bruno Roy, Arianna Rampini

P4-#4603 *Robust Reward Modeling via Causal Rubrics*
Pragya Srivastava, Harman Singh, Rahul Madhavan, Gandharv Patil, Sravanti Addepalli, Arun Suggala, Rengarajan Aravamudhan, Soumya Sharma, Anirban Laha, Aravindan Raghuveer, Karthikeyan Shanmugam, Doina Precup

P4-#4716 *When Greedy Wins: Emergent Exploitation Bias in Meta-Bandit LLM Training*
Sanxing Chen, Xiaoyin Chen, Yukun Huang, Roy Xie, Bhuwan Dhingra

P4-#5008 *µLO: Compute-Efficient Meta-Generalization of Learned Optimizers*
Benjamin Thérien, Charles-Étienne Joseph, Boris Knyazev, Edouard Oyallon, Irina Rish, Eugene Belilovsky

Mila @ ICLR 2026
24.04.2026

Follow Mila researchers at ICLR 2026
Suivez les chercheuses et chercheurs de Mila à ICLR 2026

Mila
Institut québécois d'intelligence artificielle
Quebec Artificial Intelligence Institute

Poster Session 3 - 10.30 a.m. (BRT)

P3-#1109 FALCON: Few-step Accurate Likelihoods for Continuous Flows
Danyal Rehman, Tara Akhound-Sadegh, Artem Gazizov, Yoshua Bengio, Alexander Tong

P4-#4414 Visual symbolic mechanisms: Emergent symbol processing in Vision Language Models
Rim Assouel, Declan Campbell, Yoshua Bengio, Taylor Webb

P3-#521 *Beyond Multi-Token Prediction: Pretraining LLMs with Future Summaries*
Divyat Mahajan, Sachin Goyal, Badr Youbi Idrissi, Mohammad Pezeshki, Ioannis Mitliagkas, David Lopez-Paz, Kartik Ahuja

P3-#802 *TGM: A Modular and Efficient Library for Machine Learning on Temporal Graphs*
Jacob Chmura, Shenyang(Andy) Huang, Tran Gia Bao Ngo, Ali Parviz, Farimah Poursafaei, Jure Leskovec, Michael Bronstein, Guillaume Rabusseau, Matthias Fey, Reihaneh Rabbany

P3-#1011 *Spinning Straw into Gold: Relabeling LLM Agent Trajectories in Hindsight for Successful Demonstrations*
Zichao Li, Gang Wu, Zichao Wang, Vlad Morariu, Ruiyi Zhang, Wanrong Zhu, Ryan Rossi, Jihyung Kil

P3-#1119 *Fast Proteome-Scale Protein Interaction Retrieval via Residue-Level Factorization*
Jianan Zhao, Zhihao Zhan, Narendra Chaudhary, Xinyu Yuan, Zuobai Zhang, Qian Cong, Jian Zhou, Sanchit Misra, Jian Tang

P3-#1612 *Building spatial world models from sparse transitional episodic memories*
Zizhan He, Maxime Daigle, Pouya Bashivan

P3-#1626 *DRBench: A Realistic Benchmark for Enterprise Deep Research*
Amirhossein Abaskohi, Tianyi Chen, Miguel Muñoz-Mármol, Curtis Fox, Amrutha Varshini Ramesh, Étienne Marcotte, Xing Han Lu, Nicolas Chapados, Spandana Gella, Christopher Pal, Alexandre Drouin, Issam dji

P3-#1824 *Generative Adversarial Post-Training Mitigates Reward Hacking in Live Human-AI Music Interaction*
Yusong Wu, Stephen Brade, Teng Ma, Tia-Jane Fowler, Enning Yang, Berker Banar, Aaron Courville, Natasha Jaques, Anna Huang

P3-#2006 *Benefits and Limitations of Communication in Multi-Agent Reasoning*
Michael Rizvi-Martel, Satwik Bhattamishra, Neil Rathi, Guillaume Rabusseau, Michael Hahn

P4-#3306 *Kaleidoscope: In-language Exams for Massively Multilingual Vision Evaluation*
Israfel Salazar, Manuel Fernández Burda, Shayekh Islam, Arshia Soltani Moakhar, Shivalika Singh, Fabian Farestam, Angelika Romanou, Danylo Boiko, Dipika Khullar, Mike Zhang, Dominik Krzemiński, Jekaterina Novikova, a Shimabucoro, Joseph Marvin Imperial, Rishabh Maheshwary, Sharad Duwal, Alfonso Amayuelas, Swati Rajwal, Jebish Purbey, Ahmed Ruby, Nicholas Popovič, Marek Suppa, Azmine Toushik Wasi, Ram Mohan Rao Kadiyala, Olga Tsymboi, Maksim Kostritsya, Bardia moakhar, Gabriel da Costa Merlin, Otávio Coletti, Maral Jabbarishiviari, MOHAMMADAMIN FARAHANIFARD, Silvia Fernandez, María Grandury, Dmitry Abulkhanov, Drishti Sharma, Andre Guarnier De Mitri, Leticia Marchezi, Setayesh Heydari, Johan S Obando Ceron, Nazar Kohut, Beyza Ermis, Desmond Elliott, Enzo Ferrante, Sara Hooker, Marzieh Fadaee

P4-#4918 *The Markovian Thinker: Architecture-Agnostic Linear Scaling of Reasoning*
Milad Aghajohari, Kamran Chitsaz, Amirhossein Kazemnejad, Sarath Chandar, Alessandro Sordoni, Aaron Courville, Siva Reddy

Poster Session 4 - 3.15 p.m. (BRT)

P3-#211 *Amortized Inference of Causal Models via Conditional Fixed-Point Iterations*
Divyat Mahajan, Jannes Gladrow, Agrin Hilmkil, Cheng Zhang, Meyer Scetbon

P3-#424 *Dual Optimistic Ascent (PI Control) is the Augmented Lagrangian Method in Disguise*
Juan Ramirez, Simon Lacoste-Julien

P3-#604 *MesaNet: Sequence Modeling by Locally Optimal Test-Time Training*
Chuqin Geng, Ziyu Zhao, Zhaoyue Wang, Haolin Ye, Yuhe Jiang, Xujie Si

P3-#810 *LogicXGNN: Grounded Logical Rules for Explaining Graph Neural Networks*
Seanie Lee, Minsu Kim, Lynn Cherif, David Dobre, Juho Lee, Sung Ju Hwang, Kenji Kawaguchi, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, Moksh J. Jain

P3-#1005 *Bridging Explainability and Embeddings: BEE Aware of Spuriousness*
Cristian D. Paduraru, Antonio Barbalau, Radu Filipescu, Andrei Nicolicioiu, Elena Burceanu

P3-#1106 *Property-Driven Protein Inverse Folding with Multi-Objective Preference Alignment*
Junqi Liu, Xiaoyang Hou, Chence Shi, Xin Liu, Zhi Yang, Jian Tang

P3-#1110 *h-MINT: Modeling Pocket-Ligand Binding with Hierarchical Molecular Interaction Network*
Yanru Qu, Yijie Zhang, Wenjuan Tan, Xiangzhe Kong, Xiangxin Zhou, Chaoran Cheng, Mathieu Blanchette, Jiaxuan You, Ge Liu

P3-#1117 *OXtal: An All-Atom Diffusion Model for Organic Crystal Structure Prediction*
Emily Jin, Andrei Nica, Kin Long Kelvin Lee, Joey Bose, Mikhail Galkin, Santiago Miret, Jarrid Rector-Brooks, Alexander Tong, Michael Bronstein, Frances Arnold, Chenghao Liu

P3-#1125 *Towards Sustainable Investment Policies Informed by Opponent Shaping*
Juan Duque, Razvan Ciuca, Ayoub Echchahed, Hugo Larochelle, Aaron Courville

P3-#1513 *A Balanced Neuro-Symbolic Approach for Commonsense Abductive Logic*
Joseph Cotnareanu, Didier Chételat, Yingxue Zhang, Mark Coates

P3-#1916 *La-Proteina: Atomistic Protein Generation via Partially Latent Flow Matching*
Tomas Geffner, Kieran Didi, Zhonglin Cao, Danny Reidenbach, Zuobai Zhang, Christian Dallago, Emine Kucukbenli, Karsten Kreis, Arash Vahdat

P4-#4502 *ADM-v2: Pursuing Full-Horizon Roll-out in Dynamics Models for Offline Policy Learning and Evaluation*
Haoxin Lin, Siyuan Xiao, Yi-Chen Li, Zhilong Zhang, Yihao Sun, Chengxing Jia, Yang Yu

P4-#4911 *Discrete Compositional Generation via General Soft Operators and Robust Reinforcement Learning*
Marco Jiralerspong, Esther Derman, Danilo Vucetic, Nikolay Malkin, Bilun Sun, Tianyu Zhang, Pierre-Luc Bacon, Gauthier Gidel

Oral Session 4 - 3.15 p.m. (BRT)

Oral 4C Visual symbolic mechanisms: Emergent symbol processing in Vision Language Models
Rim Assouel, Declan Campbell, Yoshua Bengio, Taylor Webb

Mila @ ICLR 2026
25.04.2026

Follow Mila researchers at ICLR 2026
Suivez les chercheuses et chercheurs de Mila à ICLR 2026

Mila
Institut québécois d'intelligence artificielle
Quebec Artificial Intelligence Institute

Poster Session 5 - 10.30 a.m. (BRT)

P3-#1718 Scaling Atomistic Protein Binder Design with Generative Pretraining and Test-Time Compute
Kieran Didi, Zuobai Zhang, Guoqing Zhou, Danny Reidenbach, Zhonglin Cao, Sooyoung Cha, Tomas Geffner, Christian Dallago, Jian Tang, Michael Bronstein, Martin Steinegger, Emine Kucukbenli, Arash Vahdat, Karsten Kreis

P3-#112 *Self-Supervised Learning from Structural Invariance*
Yipeng Zhang, Hafez Ghaemi, Jungyoon Lee, Shahab Bakhtiari, Eilif B Muller, Laurent Charlin

P3-#202 *Latent Veracity Inference for Identifying Errors in Stepwise Reasoning*
Minsu Kim, Jean-Pierre Falet, Oliver Richardson, Xiaoyin Chen, Moksh Jain, Sungjin Ahn, Sungsoo Ahn, Yoshua Bengio

P3-#423 *Towards Learned Optimization Free Lunch*
Abhinav Moudgil, Boris Knyazev, Eugene Belilovsky

P3-#504 *Egalitarian Gradient Descent: A Simple Approach to Accelerated Grokking*
Ali Saheb Pasand, Elvis Dohmatob

P3-#825 *Generalised Flow Maps for Few-Step Generative Modelling on Riemannian Manifolds*
Oscar Davis, Nicholas Boffi, Michael Albergo, Michael Bronstein, Joey Bose

P3-#1014 *Scaling Laws and Symmetry, Evidence from Neural Force Fields*
Nhat Khang Ngo, Siamak Ravanbakhsh

P3-#1501 *Embedding-Based Context-Aware Reranker*
Ye Yuan, Mohammad Amin Shabani, Siqi Liu

P3-#1521 *Learning From the Past with Cascading Eligibility Traces*
Tokiniaina Raharison Ralambomihanta, Ivan Anokhin, Roman Pogodin, Samira Ebrahimi Kahou, Jonathan Cornford, Blake A Richards

P4-#3902 *On Fairness of Task Arithmetic: The Role of Task Vectors*
Laura Gomezjurado Gonzalez , Hiroki Naganuma, Kotaro Yoshida, Takafumi Horie, Yuji Naraki, Ryotaro Shimizu

P4-#4512 *Self-Predictive Representations for Combinatorial Generalization in Behavioral Cloning*
Daniel Lawson, Adriana Hugessen, Charlotte Cloutier, Glen Berseth, Khimya Khetarpal

P4-#4608 *Temporal Representations for Exploration: Learning Complex Exploratory Behavior without Extrinsic Rewards*
Faisal Mohamed, Catherine Ji, Benjamin Eysenbach, Glen Berseth

P4-#4613 *Relative Entropy Pathwise Policy Optimization*
Claas Voelcker, Axel Brunnbauer, Marcel Hussing, Michal Nauman, Pieter Abbeel, Radu Grosu, Eric Eaton, Amir-massoud Farahmand, Igor Gilitschenski

P4-#5218 *The Geometry and Topology of Representations: the Manifolds of Modular Addition*
Gabriela Moisescu-Pareja, Gavin McCracken, Harley Wiltzer, Colin Daniels, Vincent Létourneau, Doina Precup, Jonathan Love

P4-#5317 *SynCoGen: Synthesizable 3D Molecule Generation via Joint Reaction and Coordinate Modeling*
Andrei Rekesh, Miruna Cretu, Dmytro Shevchuk, Pietro Lio, Robert Batey, Mike Tyers, Michał Koziarski, Chenghao Liu

Oral Session 5 - 10.30 a.m. (BRT)

Oral 5A Planner Aware Path Learning in Diffusion Language Models Training
Zhangzhi Peng, Zachary Bezemek, Jarrid Rector-Brooks, Shuibai Zhang, Michael Bronstein, Anru Zhang, Joey Bose, Alexander Tong

Poster Session 6 - 3.15 p.m. (BRT)

P3-#614 Planner Aware Path Learning in Diffusion Language Models Training
Zhangzhi Peng, Zachary Bezemek, Jarrid Rector-Brooks, Shuibai Zhang, Michael Bronstein, Anru Zhang, Joey Bose, Alexander Tong

P3-#222 *RAEE: A Robust Retrieval-Augmented Early Exit Framework for Efficient Inference*
LIANMING HUANG, Shangyu Wu, Yufei CUI, Ying Xiong, Haibo Hu, Xue Liu, Tei-Wei Kuo, Nan Guan, Chun Jason Xue

P3-#603 *Diffusion Alignment as Variataional Expectation-Maximization*
Jaewoo Lee, Minsu Kim, Sanghyeok Choi, Inhyuck Song, Sujin Yun, Hyeongyu Kang, Woocheol Shin, Taeyoung Yun, Kiyoung Om, Jinkyoo Park

P3-#1816 *Why Less is More (Sometimes): A Theory of Data Curation*
Elvis Dohmatob, Mohammad Pezeshki, Reyhane Askari Hemmat

P4-#4409 *The Expressive Limits of Diagonal SSMs for State-Tracking*
Mehran Shakerinava, Behnoush Khavari, Siamak Ravanbakhsh, Sarath Chandar

P4-#4410 *Defining and quantifying compositional structure*
Eric Elmoznino, Guillaume Lajoie

P4-#4515 *Simplicial Embeddings Improve Sample Efficiency in Actor-Critic Agents*
Johan S Obando Ceron, Walter Mayor, Samuel Lavoie, Scott Fujimoto, Aaron Courville, Pablo Samuel Castro

Oral Session 6 - 3.15 p.m. (BRT)

Oral 6F Scaling Atomistic Protein Binder Design with Generative Pretraining and Test-Time Compute
Kieran Didi, Zuobai Zhang, Guoqing Zhou, Danny Reidenbach, Zhonglin Cao, Sooyoung Cha, Tomas Geffner, Christian Dallago, Jian Tang, Michael Bronstein, Martin Steinegger, Emine Kucukbenli, Arash Vahdat, Karsten Kreis

Oral 6F FALCON: Few-step Accurate Likelihoods for Continuous Flows
Danyal Rehman, Tara Akhound-Sadegh, Artem Gazizov, Yoshua Bengio, Alexander Tong