Mila @ NeurIPS 2024
12.10.2024

Follow Mila researchers at NeurIPS 2024
Suivez les chercheur·se·s de Mila à NeurIPS 2024

Mila's booth / kiosque de Mila
#104 12 PM - 5 PM PST

12:15 PM Akshatha Arodi - *CableInspect-AD: An Expert-Annotated Anomaly Detection Dataset*

2:30 PM Raymond Chua - *Learning Successor Features the Simple Way*

3:15 PM Oscar Mañas - *Controlling Multimodal LLMs via Reward-guided Decoding*

4:15 PM Le Zhang - *Assessing and Learning Alignment of Unimodal Vision and Language Models*

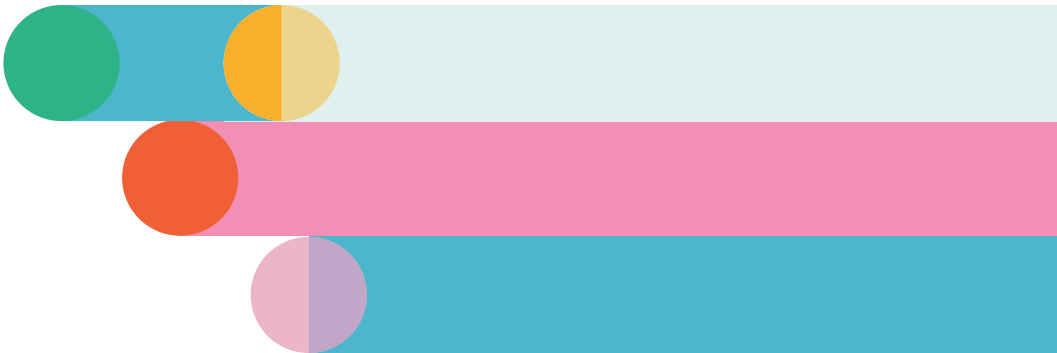

Mila @ NeurIPS 2024
12.11.2024

Follow Mila researchers at NeurIPS 2024
Suivez les chercheur·se·s de Mila à NeurIPS 2024

Mila's Booth / kiosque de Mila
#104 10 AM - 5 PM PST

10.15 AM Moksh Jain - *Amortizing intractable inference in diffusion models for vision, language, and control*

11.15 AM Aniket Didolkar - *Metacognitive Capabilities of LLMs: An Exploration in Mathematical Problem Solving*

12.15 PM Shahrad Mohammadzadeh, Juan David Guerra - *Hallucination Detox: Sensitive Neuron Dropout (SeND) for Large Language Model Training and Epistemic Integrity of Large Language Models*

2.30 PM Francesco Paissan - *Listenable Maps for Zero-Shot Audio Classifiers*

3.15 PM Benno Krojer - *Learning Action and Reasoning-Centric Image Editing from Videos and Simulation*

4.15 PM Daniel Levy - *SymmCD: Symmetry-Preserving Crystal Generation with Diffusion Models*

Poster Session 1 -
11 AM - 2 PM PST

#1110 - *How Molecules Impact Cells: Unlocking Contrastive PhenoMolecular Retrieval* - Philip Fradkin, Puria Azadi Moghadam, Karush Suri, Frederik Wenkel, Ali Bashashati, Maciej Sypetkowski, Dominique Beaini

#1606 - *Learning Action and Reasoning-Centric Image Editing from Videos and Simulation* - Benno Krojer, Dheeraj Vattikonda, Luis Lara, Varun Jampani, Eva Portelance, Christopher Pal, Siva Reddy

#2403 - *Stress-Testing Capability Elicitation With Password-Locked Models* - Ryan Greenblatt, Fabien Roger, Dmitrii Krasheninnikov, David Krueger

#2510 - *Self-Consuming Generative Models with Curated Data Provably Optimize Human Preferences* - Damien Ferbach, Quentin Bertrand, Joey Bose, Gauthier Gidel

#2509 *ET-Flow: Equivariant Flow-Matching for Molecular Conformer Generation* - Majdi Hassan, Nikhil Shenoy, Jungyoon Lee, Hannes Stärk, Stephan Thaler, Dominique Beaini

#3506 *Density-based User Representation using Gaussian Process Regression for Multi-interest Personalized Retrieval* - Haolun Wu, Ofer Meshi, Masrour Zoghi, Fernando Diaz, Xue Liu, Craig Boutilier, Maryam Karimzadehgan

#4939 *The High Line: Exact Risk and Learning Rate Curves of Stochastic Adaptive Learning Rate Algorithms* - Elizabeth Collins-Woodfin, Inbar Seroussi, Begoña García Malaxechebarría, Andrew Mackenzie, Elliot Paquette, Courtney Paquette

#5210 *CableInspect-AD: An Expert-Annotated Anomaly Detection Dataset* - Akshatha Arodi, Margaux Luck, Jean-Luc Bedwani, Aldo Zaimi, Ge Li, Nicolas Pouliot, Julien Beaudry, Gaetan Marceau Caron

#6805 *Offline Multitask Representation Learning for Reinforcement Learning* - Haque Ishfaq, Thanh Nguyen-Tang, Songtao Feng, Raman Arora, Mengdi Wang, Ming Yin, Doina Precup

Poster Session 2 -
4:30 - 7:30 PM PST

#1104 *Reactzyme: A Benchmark for Enzyme-Reaction Prediction* - Chenqing Hua, Bozitao Zhong, Sitao Luan, Liang Hong, Guy Wolf, Doina Precup, Shuangjia Zheng

#1107 *MSA Generation with Seqs2Seqs Pretraining: Advancing Protein Structure Predictions* - Le Zhang, Jiayang Chen, Tao Shen, Yu Li, Siqi Sun

#1211 *Multi-Scale Representation Learning for Protein Fitness Prediction* - Zuobai Zhang, Pascal Notin, Yining Huang, Aurelie Lozano, Vijil Chenthamarakshan, Debora Susan Marks, Payel Das, Jian Tang

#2008 *Efficient Leverage Score Sampling for Tensor Train Decomposition* - Vivek Bharadwaj, Beheshteh T. Rakhshan, Osman Asif Malik, Guillaume Rabusseau

#2606 *Doob's Lagrangian: A Sample-Efficient Variational Approach to Transition Path Sampling* - Yuanqi Du, Michael Plainer, Rob Brekelmans, Chenru Duan, Frank Noé, Carla P. Gomes, Alán Aspuru-Guzik, Kirill Neklyudov

#2705 *On improved Conditioning Mechanisms and Pre-training Strategies for Diffusion Models* - Tariq Berrada, Pietro Astolfi, Melissa Hall, Reyhane Askari Hemmat, Yohann Benchetrit, Marton Havasi, Matthew J. Muckley, Karteek Alahari, Adriana Romero-Soriano, Jakob Verbeek, Michal Drozdzal

#4801 *Normalization and effective learning rates in reinforcement learning* - Clare Lyle, Zeyu Zheng, Khimya Khetarpal, James Martens, Hado van Hasselt, Razvan Pascanu, Will Dabney

#5101 *RepLiQA: A Question-Answering Dataset for Benchmarking LLMs on Unseen Reference Content* - Joao Monteiro, Pierre-Andre Noel, Étienne Marcotte, Sai Rajeswar, Valentina Zantedeschi, David Vazquez, Nicolas Chapados, Christopher Pal, Perouz Taslakian

#5308 *WorkArena++: Towards Compositional Planning and Reasoning-based Common Knowledge Work Tasks* - Léo Boisvert, Megh Thakkar, Maxime Gasse, Massimo Caccia, Thibault Le Sellier De Chezelles, Quentin Cappart, Nicolas Chapados, Alexandre Lacoste, Alexandre Drouin

#5708 *4+3 Phases of Compute-Optimal Neural Scaling Laws* - Elliot Paquette, Courtney Paquette, Lechao Xiao, Jeffrey Pennington

#6109 *Conformal Inverse Optimization* - Bo Lin, Érick Delage, Timothy Chan

#6401 *Improving Deep Reinforcement Learning by Reducing the Chain Effect of Value and Policy Churn* - Hongyao Tang, Glen Berseth

#6910 *Balancing Context Length and Mixing Times for Reinforcement Learning at Scale* - Matthew D Riemer, Khimya Khetarpal, Janarthanan Rajendran, Sarath Chandar

#7210 *Sequence-Augmented SE(3)-Flow Matching For Conditional Protein Generation* - Guillaume Huguet, James Vuckovic, Kilian FATRAS, Eric Thibodeau-Laufer, Pablo Lemos, Riashat Islam, Cheng-Hao Liu, Jarrid Rector-Brooks, Tara Akhound-Sadegh, Michael M. Bronstein, Alexander Tong, Joey Bose

Mila @ ICML 2024
12.12.2024

Follow Mila researchers at NeurIPS 2024
Suivez les chercheur·se·s de Mila à NeurIPS 2024

Mila's Booth / kiosque de Mila
#104 10 AM - 4 PM PST

10.15 AM Jonas Ngnawé - *Detecting Brittle Decisions for Free: Leveraging Margin Consistency in Deep Robust Classifiers*

11.15 AM Prakhar Ganesh - *Different Horses for Different Courses: Comparing Bias Mitigation Algorithms in ML*

12.15 PM Bonaventure F. P. Dossou - *From Preservation to Progress: Expanding Language AI Frontiers for Low-Resource Communities*

2.30 PM Arthur Ouaknine - *Tackling Climate Change with Machine Learning*

3.15 PM Saba Ahmadi - *VisMin: Visual Minimal-Change Understanding*

Poster Session 3 - 11 AM - 2 PM PST

#1100 - *RGFN: Synthesizable Molecular Generation Using GFlowNets* - Michał Koziarski, Andrei Rekesh, Dmytro Shevchuk, Almer M. van der Sloot, Piotr Gaiński, Yoshua Bengio, Cheng-Hao Liu, Mike Tyers, Robert A. Batey

#2104 - *Reproducibility Study on Adversarial Attacks Against Robust Transformer Trackers* - Fatemeh Nourilenjan Nokabadi, Jean-Francois Lalonde, Christian Gagné

#2109 - *On the Scalability of Certified Adversarial Robustness with Generated Data* - Thomas Altstidl, David Dobre, Arthur Kosmala, Bjoern Eskofier, Gauthier Gidel, Leo Schwinn

#2411 - *Metric Flow Matching for Smooth Interpolations on the Data Manifold* - Kacper Kapusniak, Peter Potaptchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael M. Bronstein, Joey Bose, Francesco Di Giovanni

#2704 *Metacognitive Capabilities of LLMs: An Exploration in Mathematical Problem Solving* - Aniket Rajiv Didolkar, Anirudh Goyal, Nan Rosemary Ke, Siyuan Guo, Michal Valko, Timothy P Lillicrap, Danilo Jimenez Rezende, Yoshua Bengio, Michael Curtis Mozer, Sanjeev Arora

#2911 *What Is Missing For Graph Homophily? Disentangling Graph Homophily For Graph Neural Networks* - Yilun Zheng, Sitao Luan, Lihui Chen

#3011 *A Foundation Model for Zero-shot Logical Query Reasoning* - Mikhail Galkin, Jincheng Zhou, Bruno Ribeiro, Jian Tang, Zhaocheng Zhu

#3204 *The Factorization Curse: Which Tokens You Predict Underlie the Reversal Curse and More* - Ouail Kitouni, Niklas Nolte, Adina Williams, Michael Rabbat, Diane Bouchacourt, Mark Ibrahim

#3800 *Towards a "Universal Translator" for Neural Dynamics at Single-Cell* - Yizi Zhang, Yanchen Wang, Donato M. Jiménez-Benetó, Zixuan Wang, Mehdi Azabou, Blake Aaron Richards, Renee Tung, Olivier Winter, International Brain Laboratory, Eva L Dyer, Liam Paninski, Cole Lincoln Hurwitz

#4702 - *Efficient Adversarial Training in LLMs with Continuous Attacks* - Sophie Xhonneux, Alessandro Sordoni, Stephan Günnemann, Gauthier Gidel, Leo Schwinn

#6001 *HardCore Generation: Generating Hard UNSAT Problems for Data Augmentation* - Joseph Cotnareanu, Zhanguang Zhang, Hui-Ling Zhen, Yingxue Zhang, Mark Coates

#6402 *QGFN: Controllable Greediness with Action Values* - Elaine Lau, Stephen Zhewen Lu, Ling Pan, Doina Precup, Emmanuel Bengio

#6510 *Predicting Future Actions of Reinforcement Learning Agents* - Stephen Chung, Scott Niekum, David Krueger

#6704 *Foundations of Multivariate Distributional Reinforcement Learning* - Harley Wiltzer, Jesse Farebrother, Arthur Gretton, Mark Rowland

#6706 *Periodic agent-state based Q-learning for POMDPs* - Amit Sinha, Matthieu Geist, Aditya Mahajan

Poster Session 4 - 4:30 - 7:30 PM PST

#2008 *Detecting Brittle Decisions for Free: Leveraging Margin Consistency in Deep Robust Classifiers* - Jonas Ngnawe, Sabyasachi Sahoo, Yann Batiste Pequignot, Frederic Precioso, Christian Gagne

#2010 *Soft Prompt Threats: Attacking Safety Alignment and Unlearning in Open-Source LLMs through the Embedding Space* - Leo Schwinn, David Dobre, Sophie Xhonneux, Gauthier Gidel, Stephan Günnemann

#3607 *VisMin: Visual Minimal-Change Understanding* - Rabiul Awal, Saba Ahmadi, Le Zhang, Aishwarya Agrawal

#3705 *Geometry of naturalistic object representations in recurrent neural network models of working memory* - Xiaoxuan Lei, Takuya Ito, Pouya Bashivan

#3803 *Towards training digitally-tied analog blocks via hybrid gradient computation* - Timothy Nest, Maxence Ernoult

#4600 *Slight Corruption in Pre-training Data Makes Better Diffusion Models* - Hao Chen, Yujin Han, Diganta Misra, Xiang Li, Kai Hu, Difan Zou, Masashi Sugiyama, Jindong Wang, Bhiksha Raj

#4705 *Cell ontology guided transcriptome foundation model* - Xinyu Yuan, Zhihao Zhan, Zuobai Zhang, Manqi Zhou, Jianan Zhao, Boyu Han, Yue Li, Jian Tang

#5502 *Wasserstein Distributionally Robust Optimization through the Lens of Structural Causal Models and Individual Fairness* - Ahmad Reza Ehyaei, Golnoosh Farnadi, Samira Samadi

#6404 *CALE: Continuous Arcade Learning Environment* - Jesse Farebrother, Pablo Samuel Castro

#6504 *Learning Successor Features the Simple Way* - Raymond Chua, Arna Ghosh, Christos Kaplanis, Blake Aaron Richards, Doina Precup

#6508 *Using Unity to Help Solve Reinforcement Learning* - Connor Brennan, Andrew Robert Williams, Omar G. Younis, Vedant Vyas, Daria Yasafova, Irina Rish

#6606 *Adaptive Exploration for Data-Efficient General Value Function Evaluations* - Arushi Jain, Josiah P. Hanna, Doina Precup

#6612 *Simplifying Constraint Inference with Inverse Reinforcement Learning* - Adriana Hugessen, Harley Wiltzer, Glen Berseth

#7101 *Amortizing intractable inference in diffusion models for vision, language, and control* - Siddarth Venkatraman, Moksh J. Jain, Luca Scimeca, Minsu Kim, Marcin Sendera, Mohsin Hasan, Luke Rowe, Sarthak Mittal, Pablo Lemos, Emmanuel Bengio, Alexandre Adam, Jarrid Rector-Brooks, Yoshua Bengio, Glen Berseth, Nikolay Malkin

Mila @ NeurIPS 2024
12.13.2024

Follow Mila researchers at NeurIPS 2024
Suivez les chercheur·se·s de Mila à NeurIPS 2024

Poster Session 5 - 11 AM - 2 PM PST

#1002 - *Trajectory Flow Matching with Applications to Clinical Time Series Modelling* - Xi Zhang, Yuan Pu, Yuki Kawamura, Andrew Loza, Yoshua Bengio, Dennis Shung, Alexander Tong

#2201 - *Harnessing small projectors and multiple views for efficient vision pretraining* - Arna Ghosh, Kumar Krishna Agrawal, Shagun Sodhani, Adam Oberman, Blake Aaron Richards

#2211 - *Expecting The Unexpected: Towards Broad Out-Of-Distribution Detection* - Charles Guille-Escuret, Pierre-Andre Noel, Ioannis Mitliagkas, David Vazquez, Joao Monteiro

#2606 - *Fisher Flow Matching for Generative Modelling over Discrete Data* - Oscar Davis, Samuel Kessler, Mircea Petrache, Ismail Ilkay Ceylan, Michael Bronstein, Joey Bose

#3103 *On the Scalability of GNNs for Molecular Graphs* - Maciej Sypetkowski, Frederik Wenkel, Farimah Poursafaei, Nia Dickson, Karush Suri, Philip Fradkin, Dominique Beaini

#3202 *Listenable Maps for Zero-Shot Audio Classifiers* - Francesco Paissan, Luca Della Libera, Mirco Ravanelli, Cem Subakan

#3301 *Interpreting Learned Feedback Patterns in Large Language Models* - Luke Marks, Amir Abdullah, Clement Neo, Rauno Arike, David Krueger, Philip Torr, Fazl Barez

#3303 *Improving Context-Aware Preference Modeling for Language Models* - Silviu Pitis, Ziang Xiao, Nicolas Le Roux, Alessandro Sordoni

#3307 *Do LLMs Build World Representations? Probing Through the Lens of State Abstraction* - Zichao Li, Yanshuai Cao, Jackie C. K. Cheung

#3606 *Grounding Multimodal Large Language Models in Actions* - Andrew Szot, Bogdan Mazoure, Harsh Agrawal, Devon Hjelm, Zsolt Kira, Alexander T Toshev

#6304 *Parseval Regularization for Continual Reinforcement Learning* - Wesley Chung, Lynn Cherif, Doina Precup, David Meger

#6410 *Action Gaps and Advantages in Continuous-Time Distributional Reinforcement Learning* - Harley Wiltzer, Marc G. Bellemare, David Meger, Patrick Shafto, Yash Jhaveri

Poster Session 6 - 4:30 - 7:30 PM PST

#1005 *Code Repair with LLMs gives an Exploration-Exploitation Tradeoff* - Hao Tang, Keya Hu, Jin Peng Zhou, Si Cheng Zhong, Wei-Long Zheng, Xujie Si, Kevin Ellis

#2002 *When is an Embedding Model More Promising than Another?* - Maxime Darrin, Philippe Formont, Ismail Ben Ayed, Jackie C. K. Cheung, Pablo Piantanida

#2600 *Improved off-policy training of diffusion samplers* - Marcin Sendera, Minsu Kim, Sarthak Mittal, Pablo Lemos, Luca Scimeca, Jarrid Rector-Brooks, Alexandre Adam, Yoshua Bengio, Nikolay Malkin

#2701 *Many-Shot In-Context Learning* - Rishabh Agarwal, Avi Singh, Lei M Zhang, Bernd Bohnet, Luis Rosias, Stephanie C.Y. Chan, Biao Zhang, Ankesh Anand, Zaheer Abbas, Azade Nova, John D Co-Reyes, Eric Chu, Feryal Behbahani, Aleksandra Faust, Hugo Larochelle

#3710 *A Generative Model of Symmetry Transformations* - James Urquhart Allingham, Bruno Mlodozeniec, Shreyas Padhy, Javier Antoran, David Krueger, Richard E. Turner, Eric Nalisnick, José Miguel Hernández-Lobato

#6110 *GenRL: Multimodal-foundation world models for generalization in embodied agents* - Pietro Mazzaglia, Tim Verbelen, Bart Dhoedt, Aaron Courville, Sai Rajeswar

#6805 *Efficient Reinforcement Learning by Discovering Neural Pathways* - Samin Yeasar Arnob, Riyasat Ohib, Sergey Plis, Amy Zhang, Alessandro Sordoni, Doina Precup