Bogdan Mazoure

Curriculum Vitae

Cambridge, United Kingdom
+44 7890 648071
bogdan.mazoure@gmail.com

Education

2019–2022 **PhD in Computer Science**, *McGill University/ MILA*, Montreal.

2017–2018 **Master's in Statistics**, *McGill University*, Montreal, *GPA – 3.95/4.0*.

2014–2017 **B.Sc. in Statistics and Computer Science**, *McGill University*, Montreal, *GPA – 3.77/4.0*.

Job Experience

2022-Now **Research Scientist**, *Apple, Machine Learning Research*.

2022 **Research Intern**, *DeepMind, Montreal*.

2022 **Student Researcher**, *Google Brain, Mountain View*.
- Pre-training state abstractions based on density models for sample-efficient finetuning.

2022 **Course lecturer**, *McGill University*.
- COMP 424 - Artificial Intelligence [link].

2021 **Course lecturer**, *UQAM University*.
- BIF7105 - Statistical methods in bioinformatics [link].

2021 **Research Intern**, *Google Brain, Mountain View*.
- Improving zero-shot generalization capabilities of offline RL methods via contrastive losses.

2020–2021 **Researcher (part-time)**, *Microsoft Redmond*.

2020 **Research Intern**, *Microsoft Research, New York*.
- Long-term rewards optimization in batch RL via short-horizon policy iteration.

2019 **Research Intern**, *Microsoft Research, Montreal*.
- Learning predictive state representations for high-dimensional sequential decision making tasks.

2018 **Research Intern**, *Nuance*.
- Improving upon current state-of-the-art architectures for ambient speech recognition.

2018–2019 **Consultant**, *Correlation One*.
- Designing technical interview questions on fundamental machine learning and probability topics used by top hedge funds (e.g. Two Sigma, Citadel).

2016–2020 **Teaching Assistant**, *McGill University*.
- Holding weekly office hours and tutorials, designing, invigilating and grading midterms and final exams.
- **Courses**: Artificial Intelligence, Introduction to Programming, Probability, Statistics, Algebra 2.

2015–2017 **Research Assistant**, *UQAM and McGill University*.
- Created a sequence alignment method combining Monte Carlo and travelling salesman problem.
- Developed a novel approach based on iterative proportional fitting to remove bias from drug discovery experiments.

Community service

2024 Organizer of "Training Agents with Foundation Models" workshop (RLC 2024)

2024 Organizer of "Generative Models for Decision-Making" workshop (ICLR 2024)

2024 Reviewer for AAAI 2022, AAAI 2021, AISTATS 2021, PLOS One, NeurIPS 2023, CoRL 2024

2021 Invited talk at "Real world RL" workshop (ICML 2021)

2021 Organizer of "Theory and foundations of continual learning" workshop (ICML 2021)

2021 Organizer of "Self-supervision for reinforcement learning" workshop (ICLR 2021)

2021 Invited talk at CAMDEA Toronto digital forum

2021 Invited talk on representation learning for reinforcement learning (UQAM University)

2019 Invited talk on transfer learning in deep learning (McGill University)

Awards

2019-2023 FRQNT doctoral stipend (Accepted for 4 years)

2016–2017 NSERC undergraduate research stipend (Accepted 2 years in a row)

Selected Publications

Conferences and Workshops

Published Szot, A.; **Mazoure, B.**; Agarwal, H., Hjelm, D.; Zsolt, K.; Toshev, A. Grounding Multimodal Large Language Models in Actions. *NeurIPS 2024*. [link]

Published Szot, A.; Schwarzer, M.; Agarwal, H., **Mazoure, B.**; Talbott, W.; Metcalf, K;, Mackraz, N.; Hjelm, D.; Toshev, A. Large Language Models as Generalizable Policies for Embodied Tasks. *ICLR 2024*. [link]

Published **Mazoure, B.**; Eysenbach, B.; Nachum, O.; Tompson, J. Contrastive Value Learning: Implicit Models for Simple Offline RL. *CoRL 2023*. [link]

Published **Mazoure, B.**; Bruce, J.; Precup, D.; Fergus, R.; Anand, A. Accelerating exploration and representation learning with offline pre-training. *ICML 2023 (ILHF workshop)*. [link]

Published Bruce, J.; Anand, A.; **Mazoure, B.**; Fergus, R. Learning About Progress From Experts. *ICLR 2023*. [link]

Published **Mazoure, B.**; Kostrikov, I.; Nachum, O.; Tompson, J. Improving Zero-shot Generalization in Offline Reinforcement Learning using Generalized Value Functions. *NeurIPS 2022*. [link]

Published **Mazoure, B.**; Mineiro, P.; Srinath, P.; Sedeh, R.S.; Precup, D.; Swaminathan, A. Improving Long-Term Metrics in Recommendation Systems using Short-Horizon Offline RL. *ECML-PKDD*, 2022. [link]

Published Li, T.; **Mazoure, B.**; Rabusseau, G. Sequential Density Estimation via Nonlinear Continuous Weighted Finite Automata. *LearnAut*, 2022.

Published **Mazoure, B.***; Ahmed, A. M.*; MacAlpine, P.; Hjelm, R D.; Kolobov, A. Cross-Trajectory Representation Learning for Zero-Shot Generalization in RL. *ICLR*, 2022. [link]

Published Doan, T; Bennani, M.; **Mazoure, B.**; Rabusseau, G.; Alquier, P. A Theoretical Analysis of Catastrophic Forgetting through the NTK Overlap Matrix. *AISTATS*, 2021. [link]

Published **Mazoure, B.***; Tachet des Combes, R.*; Doan, T; Bachman, P.; Hjelm, D. Deep Reinforcement and InfoMax Learning. *NeurIPS*, 2020. [link]

Published **Mazoure, B.***; Li, T.*; Precup, D.; Rabusseau, G. Efficient Planning under Partial Observability with Unnormalized Q Functions and Spectral Learning. *AISTATS* , 2020. [link]

Published **Mazoure, B.***; Doan, T.*; Durand, A.; Hjelm, R D.; Pineau, J. Leveraging exploration in off-policy algorithms via normalizing flows. *CoRL (Spotlight)*, 2019. [link]

Published Wang, J.; **Mazoure, B.**; McCracken, G.; Venuto, D.; Durand, A. PAC-Bayesian Analysis of Counterfactual Risk in Stochastic Contextual Bandits. *RLDM*, 2019. [link]

Published Norouzian, A.; **Mazoure, B.**; Connolly, D.; Willett, D. Classifying audio segments into system-directed and non-system-directed. *ICASSP*, 2019. [link]

Published Doan, T.; Monteiro, J.; Albuquerque, I.; **Mazoure, B.**; Durand, A.; Pineau, J. ; Hjelm, R D. Online Adaptative Curriculum Learning for GANs. *AAAI (Spotlight)*, 2019. [link]

Published **Mazoure, B.**; Doan, T. Learning emojis distributions with a generative model. *EMNLP Workshop on Subjectivity*, 2018. [link]

Journals

Published **Mazoure, B.***; Doan, T.*; Li, T; Makarenkov, V.; Pineau, J.; Precup, D.; Rabusseau, G. Representation of Reinforcement Learning Policies in Reproducing Kernel Hilbert Spaces. *Journal of Artificial Intelligence Research*, 2022. [link]

Published **Mazoure, B.***; Makarenkov, V.*; Rabusseau, G.; Legendre, P. Horizontal gene transfer and recombination analysis of SARS-CoV-2 genes helps discover its close relatives and shed light on its origin. *BMC Ecology and Evolution*, 2021.

Published Mazza-Anthony, C.; **Mazoure, B.**; Coates, M. Learning Gaussian Graphical Models with Ordered Weighted L_1 Regularization. *IEEE Transactions on Signal Processing*, 2020. [link]

Published **Mazoure, B.**; Nadon, R.; Makarenkov, V. Identification and correction of spatial bias are essential for obtaining quality data in high-throughput screening technologies. *Nature Scientific Reports*, 2017. [link]

Published **Mazoure, B.**; Caraus, I.; Nadon, R.; Makarenkov, V. Identification and correction of additive and multiplicative spatial biases in experimental high-throughput screening. *SLAS Discovery*, 2017. [link]

Published Caraus, I.; **Mazoure, B.**; Nadon, R.; Makarenkov, V. Detecting and removing multiplicative spatial bias in high-throughput screening technologies. *Bioinformatics*, 2017. [link]

Under submission

Submitted Klissarov, M.; Hjelm, D.; Toshev, A. **Mazoure, B.** On the Modeling Capabilities of Large Language Models for Sequential Decision Making. [link]

Submitted **Mazoure, B.**; Talbott, W.; Bautista, M.A.; Hjelm, D.; Toshev, A; Susskind, J. Value function estimation using conditional diffusion models for control. [link]

Submitted Hjelm, R D.*; **Mazoure, B.***; Golemo, F.; Frujeri, F.; Jalobeanu, M.; Kolobov, A. The Sandbox Environment for Generalizable Agent Research (SEGAR). [link]

Relevant Skills

Intermediate Java, Git, LaTeX, Unix, Windows, Mathematica, C, Matlab, Flask, Azure/ AWS/ Google computing, Speech recognition

Advanced Statistical and probability theory, Linear algebra, Reinforcement learning, Representation learning, Large Language Models, Diffusion models, R, Python (Jax, PyTorch, TensorFlow), Algorithms and data structures, Distributed computing, W&B (experiment tracking).