

AVISHEK (JOEY) BOSE

joeey.bose@cs.ox.ac.uk

EDUCATION

University of Oxford
Post-Doc

February 2024 - Present
Supervisor: **Prof. Michael Bronstein**

**McGill University &
Mila - Quebec AI Institute**
Phd in Computer Science

September 2018 - September 2023
Supervisors: **Prof. Prakash Panangaden**
Prof. William Hamilton
Prof. Gauthier Gidel

University of Toronto
M.aSc Department of Electrical and Computer Engineering

September 2017 - August 2018
Supervisor: **Prof. Parham Aarabi**

University of Toronto
B.aSc Department of Electrical and Computer Engineering
Minor in Mechatronics

September 2012 - April 2017

EDUCATION

University of Oxford
Post-Doc

February 2024 - Present
Supervisor: **Prof. Michael Bronstein**

**McGill University &
Mila - Quebec AI Institute**
Phd in Computer Science

September 2018 - September 2023
Supervisors: **Prof. Prakash Panangaden**
Prof. William Hamilton
Prof. Gauthier Gidel

University of Toronto
M.aSc Department of Electrical and Computer Engineering

September 2017 - August 2018
Supervisor: **Prof. Parham Aarabi**

University of Toronto
B.aSc Department of Electrical and Computer Engineering
Minor in Mechatronics

September 2012 - April 2017

POSITIONS HELD

Mila
Mila Affiliate Faculty member

Sept 2024 - Present

- Affiliate member of Mila which formally allows for supervision of graduate students and collaboration with other academic faculty at Mila. This is equivalent to the status of Adjunct Faculty at Mila.

Dreamfold
Distinguished Machine Learning Scientist

July 2023 - Present

- Research on Equivariant Generative models for de novo proteins design.

Qualcomm AI research
Research Intern, Supervisors: Johann Bremer and Taco Cohen

Oct 2022 - Feb 2023

- Research on Equivariant Diffusion Models for Trajectory Optimization in Reinforcement Learning.

Facebook AI Research

May 2021 - Dec 2021

Research Intern, Supervisor: Aditya Grover

- Research on investigating the causal structure learned by Deep Generative Models.
- Research on extending score-based generative models to Riemannian manifolds.

D-ID

Jun 2020 - Dec 2020

Consultant

- Consultant to D-ID research interest in privacy-preserving Machine Learning based on the acquisition of FaceShield.

Uber AI

May 2019 - Aug 2019

Research Intern, Supervisors: Ankit Jain and Piero Molino

- Research on Meta-Learning for Graph Representation Learning with real-world applications to Uber Eats. Published NeurIPS 2019 workshop paper.

FaceShield (Acquired in 2020 by D-ID)

Aug 2018 - Jun 2020

Founder and CEO

- Founded FaceShield Inc. as a way to bring research done during Masters to the masses. The main mandate of FaceShield is to increase awareness of issues surrounding digital privacy.
- Created a tool that allows anyone to apply a privacy filter to their own photos.

Borealis AI

May 2017 - July 2018

Research Intern, Supervisor: Yanshuai Cao

- Published a paper to ACL on Adversarial Contrastive Estimation that effectively improves Noise Contrastive Estimation approaches by using an Adversarial Adaptive Conditional Negative Sampler to sample harder negatives leading to better and faster convergence.
- Improved the performance of word embedding, knowledge graph embeddings, and ordered embeddings on multiple benchmark tasks

Architech

May 2015 - Aug 2016

Junior Machine Learning Engineer

- Created a novel gaze-tracking algorithm using input video inputs from a webcam's
- Applied Topic Modeling techniques to Twitter to discover new emergent sentiments on news stories before they became viral

University of Toronto

May 2014 - Aug 2014

Research Assistant

- Created demos of Veillance Flux and Augmented Reality with the help of Prof. Steve Mann.
- Hardware technologies used include Meta Space Glasses, Microsoft Kinect, and Depth Sense.
- Co-author in a paper published in IEEE GEM 2014 conference

TEACHING EXPERIENCE

McGill University

Sept 2022 - Dec 2022

Co-Instructor

- Formulated a new graduate seminar course COMP760 at McGill/Mila at the intersection of geometric deep learning and generative models.

- The course was co-taught with my supervisor Prof. Prakash Panangaden and was the first ever *graduate seminar course* co-taught by a PhD student at McGill. Course Website.

McGill University
Head TA

Jan 2019 - Dec 2019

- Responsible for the successful organization of COMP 551: Applied ML Graduate Course.

University of Toronto
Co-Head TA and Co-Lecturer

Sept 2017 - Dec 2017

- Created and delivered Lectures and Tutorials for a probability course aimed at 3rd/4th year ECE students.
- Responsible for creating and marking of quizzes and assignments and overall logistics of the course.

JOURNAL AND CONFERENCE PUBLICATIONS

* denotes equal contribution, † denotes equal advising.

1. D. Ferbach, Q. Bertrand, **A.J. Bose**, G. Gidel, “Self-Consuming Generative Models with Curated Data Provably Optimize Human Preferences” Neural Information Processing Systems (NeurIPS 2024) **Spotlight**. [arXiv link](#)
2. G. Huguet*, J. Vuckovic*, K. Fatras, E Thibodeau-Laufer, P. Lemos, R. Islam, C.H. Liu, J. Rector-Brooks, T. Akhound-Sadegh, M. Bronstein, A. Tong,† **A.J. Bose**†, “Sequence-Augmented SE(3)-Flow Matching for Conditional Protein Backbone Generation” Neural Information Processing Systems (NeurIPS 2024). [arXiv link](#)
3. K. Kapusniak, P. Potaptchik, T. Reu, L. Zhang, A. Tong, M. Bronstein, **A.J. Bose**, F. Di Giovanni, “Metric Flow Matching for Smooth Interpolations on the Data Manifold” Neural Information Processing Systems (NeurIPS 2024). [arXiv link](#)
4. O. Davis, S. Kessler, M. Petrache, I.I. Ceylan, M. Bronstein, **A.J. Bose**, “Fisher Flow Matching for Generative Modeling over Discrete Data” Neural Information Processing Systems (NeurIPS 2024). [arXiv link](#)
5. T. Akhound-Sadegh*, J. Rector-Brooks*, **A.J. Bose***, S. Mittal, P. Lemos, C.H. Liu, M. Sendera, S. Ravanbaksh, G. Gidel, Y. Bengio, N. Malkin, A. Tong “Iterated Denoising Energy Matching for Sampling from Boltzmann Densities” International Conference on Machine Learning (ICML 2024). [arXiv link](#)
6. **A. J. Bose***, T. Akhound-Sadegh*, K. Fatras, G. Huguet, J. Rector-Brooks, C.H. Liu, A.C. Nica, M. Korablyov, M. Bronstein, A. Tong, “SE(3) Stochastic Flow Matching for Protein For Protein Backbone Generation”, International Conference on Learning Representations (ICLR 2024) **Spotlight**. [arXiv link](#)
7. Q. Bertrand, **A. J. Bose**, A. Duplessis, M. Jiralerspong, G. Gidel, “On the Stability of Iterative Retraining of Generative Models on their own Data”, International Conference on Learning Representations (ICLR 2024) **Spotlight**. [arXiv link](#)
8. J. Bremer*, **A.J. Bose***, P. de Haan, T. Cohen, “EDGI: Equivariant diffusion for planning with embodied agents” Neural Information Processing Systems (NeurIPS 2023) [arXiv link](#).
9. M. Jiralerspong, **A.J. Bose**, G. Gidel, “Feature Likelihood Divergence: Evaluating Generalization of Generative Models Using Samples” Neural Information Processing Systems (NeurIPS 2023). [arXiv link](#)

10. D. Ferbach*, C. Tsirigotis*, G. Gidel, **A.J. Bose**, “A General Framework For Proving The Equivariant Strong Lottery Ticket Hypothesis” International Conference on Learning Representations (ICLR) 2023 **OpenReview link**
11. **A.J. Bose**, R. P. Monti, A. Grover “Controllable Generative Modelling via Causal Reasoning” Transactions on Machine Learning Research 2022 (TMLR 2022) **OpenReview link**
12. C. Huang*, M. Aghajohari*, **A.J. Bose**, P. Panangaden, A. Courville “Riemannian Diffusion Models” Neural Information Processing Systems (NeurIPS 2022) **arXiv link**
13. H. Ben-Hamu, S. Cohen, **A.J. Bose**, B. Amos, M. Nickel, A. Grover, R. Chen, Y. Lipman “Matching Normalizing Flows and Probability Paths on Manifolds” International Conference on Machine Learning (ICML) 2022 **arXiv link**.
14. A. Mladenović*, **A.J. Bose***, H. Berard*, W.L. Hamilton, S. Lacoste-Julien, P. Vincent, G.Gidel “Online Adversarial Attacks” International Conference on Learning Representations (ICLR) 2022. **arXiv link**
15. N. Dziri, A. Madotto, O. Zaiane, **A. J. Bose** “Neural Path Hunter: Reducing Hallucination in Dialogue Systems via Path Grounding” Empirical Methods in Natural Language Processing 2021 **arXiv link**
16. **A. J. Bose***, G. Gidel*, H. Berard*, A. Cianflone, P. Vincenct, S.L. Julien, W. L. Hamilton. (2020) “Adversarial Example Games” Neural Information Processing Systems 2020 (NeurIPS) **arXiv link**.
17. K. Ahrabian*, A. Feizi*, Y. Salehi*, W. L. Hamilton, **A. J. Bose** “Structure Aware Negative Sampling in Knowledge Graphs” Empirical Methods in Natural Language Processing 2020 (EMNLP). **arXiv link**
18. **A. J. Bose**, A. Smofsky, R. Liao, P. Panangaden, W. L. Hamilton. (2019) “Latent Variable Modeling with Hyperbolic Normalizing Flows” International Conference of Machine Learning 2020. **arXiv link**
19. X. Peng, H. Saghir, J. Kang, T. Long, **A. J. Bose**, Y. Cao, J. Cheung (2019) “A Cross-Domain Transferable Neural Coherence Model.” In Proceedings of Association for Computational Linguistics 2019 (ACL). **arXiv link**
20. **A. J. Bose**, W. L. Hamilton. (2019) “Compositional Invariance Constraints for Graph Embeddings” International Conference of Machine Learning 2019 (ICML). **arXiv link**
21. A. Cianflone, Z. Ahmed, R. Islam, **A. J. Bose**, W.L. Hamilton. (2019) “Discrete off-policy policy gradient using continuous relaxations” Reinforcement Learning and Decision Making 2019 (RLDM). **paper link**
22. **A. J. Bose**, P. Aarabi (2018) “ Adversarial Attacks on Face Detectors using Neural Net based Constrained Optimization.”, IEEE MMSP, Vancouver, Canada. **ORAL and Best Paper Nominee arXiv link**
23. **A. J. Bose***, Y. Cao*, H. Ling* (2018) “ Adversarial Contrastive Estimation.” **ORAL** In Proceedings of Association for Computational Linguistics 2018, Melbourne, Australia. **arXiv link**
24. R. Janzen, S. N. Yasrebi, **A. J. Bose**, A. Subramanian, S. Mann, “ Walking through sight: Seeing the ability to see in a 3-D augmented reality environment”, Proc. IEEE Gaming Entertainment Media, pp. 313-4, 2014. **paper link**

WORKSHOP PUBLICATIONS

1. K. Panford-Quainoo, **A. J. Bose** (2019), M. Defferrard, “Bilateral Trade Modeling with Graph Neural Networks” International Conference on Learning Representations 2020 Workshop.
2. **A. J. Bose**, A. Cianflone, W. L. Hamilton. (2019) “Graph Attacks with Latent Variable Noise Modelling” NeurIPS Graph Representation Learning Workshop, Montreal, Canada.
3. **A. J. Bose**, P. Aarabi (2019) “Virtual Fakes: DeepFakes for Virtual Reality.”, IEEE MMSP, Kuala Lumpur, Malaysia, **IEEE link**
4. P.N. Ward*, A. Smofsky*, **A. J. Bose** (2019) “Improving Exploration in Soft-Actor-Critic with Normalizing Flows Policies” Invertible Neural Networks and Normalizing Flows Workshop ICML 2019. **Spotlight talk arXiv link**
5. **A. J. Bose**, W. L. Hamilton. (2018) “Compositional Fairness Constraints for Graph Embeddings” NeurIPS Relational Representation Learning Workshop, Montreal, Canada. **paper link**
6. **A. J. Bose***, H. Ling*, Y. Cao (2018) “Compositional Hard Negative Mining for Visual Semantic Embeddings via an Adversary” NeurIPS 2018 Visually Grounded Language and Interaction Workshop, Montreal, Canada. **paper link**
7. **A. J. Bose**, A. Jain, P. Molino, W. L. Hamilton. (2019) “Meta-Graph: Few Shot Link Prediction via Meta Learning”, NeurIPS 2019 Graph Representation Learning Workshop **arXiv link**.

PREPRINTS AND PATENTS

1. M. Skerta*, L. Atanackovic*, **A. J. Bose**, A. Tong, K. Neklyudov. (Preprint 2024) “The Superposition of Diffusion Models” **Openreview link**
2. J. Rector-Brooks, M. Hasan, Z. Peng, Z. Quinn, C.H. Liu, S. Mittal, N. Dziri, M. Bronstein, Y. Bengio, P. Chatterjee, A. Tong[†], **A. J. Bose**[†]. (2024) “Steering Masked Discrete Diffusion Models via Discrete Denoising Posterior Prediction” arXiv link
3. **A.J. Bose**, M. Brubaker, I. Kobyzev “Equivariant Finite Normalizing Flows” (2022) Preprint. **arXiv link**
4. **A. J. Bose**, and P. Aarabi. “Disruption of Face Detection” U.S. Patent No. P7847US00. 13 Sept. 2018.

WORKSHOPS ORGANIZED

1. **A.J. Bose**, E. Mathieu, C. LeLan, I. Chami, F. Sala, M. Nickel, C. De Sa, C. Ré, “Differential Geometry meets Deep Learning @ NeurIPS 2020” **DiffGeo4DL Website**
2. T. Akhound-Sadegh, M. Skreta, Y. Du, S. Mittal, **A.J. Bose**, A. Tong, K. Nekyludov, M. Bronstein, M. Welling, A. Doucet, A. Hyvrinen, “Frontiers in Probabilistic Inference: Sampling Meets Learning” @ ICLR 2025 **FPI Website**

ACADEMIC ACHIEVEMENTS

NSERC Post-Doctoral Fellowship

Sept 2023

- Awarded a Post-doctoral fellowship by NSERC Canada based on academic excellence in research conducted during doctoral studies. The award comes with \$140,000 over 2 years.

ICML 2024 Top Reviewer Award

July 2024

- Awarded a top reviewer for being among the top reviewers at July 2024. This award is based on the feedback of meta-reviewers and is awarded for excellence in service while providing constructive feedback to authors.

- NeurIPS 2022 Top Reviewer Award** Oct 2022
- Awarded a top reviewer for being among the top reviewers at NeurIPS 2022. This award is based on the feedback of meta-reviewers and is awarded for excellence in service while providing constructive feedback to authors.
- ICML 2020 Top Reviewer Award** Sept 2020
- Awarded a top reviewer for being among the top 33% best reviewers at ICML 2020. This award is based on the feedback of meta-reviewers and is awarded for excellence in service while providing constructive feedback to authors.
- IVADO Ph.D. Fellowship** June 2019
- The goal of the excellence fellowship program is to support promising students in their training as future highly qualified personnel (researchers, professors, professionals) and more generally, future actors in the field of data science, mainly in IVADO members areas of excellence: operations research, machine learning, decision sciences. The award comes with \$100,000 over a 4 year span.
- IEEE MMSP Best Paper Honorable Mention** Aug 2018
- Runner-up prize for the best paper at IEEE MMSP. The paper focused on crafting adversarial attacks against modern face detectors and showed the existence of such attacks using parametric methods.
- Top 5 Impressive Graduating Students** Nov 2018
- Presented to 5 students graduating from the University of Toronto who have made the most of their time. This led to an interview with Uoft News and an article feature.
- Young Stars Estes Award** May 2018
- Awarded on a competitive basis to graduate students and post-docs attending the Deep, fast, and Shallow Learning in Humans and Machines conference at Indiana State University.
- Gordon Slemon Design Award** Oct 2017
- The Gordon Slemon Design Award is awarded for excellence in engineering design. It is awarded for the best 4th year design project as judged by the department based on effective planning, scheduling, reporting, and excellence in design, execution, creativity, etc. The award is in the form of a \$1000 cash prize along with an engraved plaque of all team members.
- Centennial Thesis Award** May 2017
- This award is offered to the fourth-year student who receives the highest grade in the 4th year Design Project. One award is given for each program: electrical and computer engineering programs. Each award is in the form of a \$500 prize and an accompanying certificate.
- Dean's List** April 2015,2017
- Awarded for academic performance for having an average higher than 80%
- Uoft TrackOne Scholarship** April 2012
- Awarded \$2000 based on academic merit

INVITED TALKS AND PANELS

- OAW AI Winter School 2025** January 2025
- A tutorial on Geometric Generative Models and their applications.

- Mediterranean ML Summer School** September 2024
- An introductory lecture on diffusion, flow matching, and other simulation-free generative models.
- ICML Geometric Representation Learning and Generative Modeling** July 2024
- Talk on using Meta-Learning for Few-Shot Link Prediction in Graph Datasets.
- University College London** May 2020
- Talk on using Meta-Learning for Few-Shot Link Prediction in Graph Datasets.
- Facebook AI Research (FAIR)** February 2020
- Talk on latent variable modeling with hyperbolic flows. The research from this version of the talk was later published at ICML 2020.
- University of Toronto & Vector Institute** October 2019
- Talk on using Meta-Learning for Few-Shot Link Prediction in Graph Datasets.
- Mila Graph Representation Learning Group** February 2019
- Talk on enforcing compositional invariance constraints for Node Embeddings in graph data structures. The research from this version of the talk was submitted to ICML 2019.
- McGill Computational Linguistics Group** November 2018
- Talk on applying Adversarial Contrastive Estimation for Image Caption Retrieval. The research from this version of the talk was later published at the NeurIPS Visually Grounded Language and Interaction Workshop.
- Facebook AI Research (FAIR)** November 2018
- Talk on enforcing compositional fairness constraints for Node Embeddings in graph data structures. The research from this version of the talk was later published at the NeurIPS Relational Representation Learning Workshop.
- GeekPwn at DefCon: AI/Robotics and Cybersecurity** June 2018
- Talk on Adversarial Attacks against break face detectors to an audience of security experts.
- Huawei AI and Security Workshop** June 2018
- Presented Research on Adversarial Attacks against face detectors to Huawei and other expert researchers in the field of adversarial machine learning.
- UofT Engineering Science Math, Statistics, and Finance Panel Discussion** April 2018
- Participated as a panelist for 2nd and 3rd year Engineering Science students. Provided advice and direction to students entering machine learning and research in general.

PRESS COVERAGE

University of Toronto News “Meet five impressive graduating students who got the most of their U of T experience” [article link](#)

University of Toronto News “U of T Engineering AI researchers design privacy filter for your photos that disables facial recognition systems” [article link](#)

Forbes “AI Researchers Create ‘Privacy Filter’ That Disrupts Facial Recognition Technology” **article link**

CBC “U of T researchers developing a tool to jam facial recognition software” **article link**

VentureBeat “University of Toronto researchers develop AI that can defeat facial recognition systems” **article link**

Toronto Star “U of T researchers design algorithm that dupes facial recognition detectors” **article link**

Science Daily “AI researchers design ‘privacy filter’ for your photos” **article link**

New York Times “When A.I.s Output Is a Threat to A.I. Itself” **article link**

COMMUNITY SERVICE

Tutorials Organized

Geometric Generative Models at Learning on Graphs conference (LoG 2024)

Reviewing

Member of the Program Committee for **ICML 2020,2021,2023**, **NeurIPS 2020,2021,2022, 2023, 2024**, **AAAI 2021**, **ICLR 2022,2023,2025**, **AISTATS 2022**, Transactions of Machine Learning Research **TMLR**, and reviewer for Transactions in Pattern Analysis and Machine Intelligence (**TPAMI**).