

# Joseph Cotnareanu

PHD STUDENT · MCGILL UNIVERSITY - ILLS, MILA

✉ joseph.cotnareanu@mail.mcgill.ca | Languages: English • French

## Education

### McGill University

Montreal, CAN

#### PHD ELECTRICAL AND COMPUTER ENGINEERING

Sept. 2024 - present

- Advisor: Dr. Mark Coates
- Topics: Abductive Logical Reasoning, Probabilistic Logical Reasoning, Reasoning with Large Language Models

### McGill University

Montreal, CAN

#### MSC ELECTRICAL AND COMPUTER ENGINEERING

Sept. 2022 - Sept. 2024

- Advisor: Dr. Mark Coates
- Thesis: "Generating Hard and Realistic Boolean (Un)Satisfiability Problems"

### McGill University

Montreal, CAN

#### B.ENG ELECTRICAL AND COMPUTER ENGINEERING

Sept. 2018 - May 2022

## Research Experience

### Graduate Student Researcher - McGill University

Montreal, CAN

#### ADVISOR: DR. MARK COATES

Aug. 2022 - Present

- "Abductive Reasoning with Probabilistic Commonsense" **[Accepted at ICML '26]**: propose a probabilistic interpretation for the uncertainty of the truth of a statement given human ambiguity and non-consensus; propose accordingly a greedy search algorithm which approximates the outcome of an MCTS search over the "tree of thoughts" of an LLM's verbalized reasoning trace for the final truth decision.
- "A Balanced Neuro-Symbolic Approach for Commonsense Abductive Logic" **[Accepted at ICLR '26]**: designing logical reasoning algorithms which leverage formal-logical, probabilistic, and machine-learning tools. By merging the three into a single hybrid approach, we achieve robustness to low-quality user input while still leveraging symbolic and interpretable reasoning strategies.
- "HardCore Generation: Generating Hard UNSAT Problems for Data Augmentation" **[Accepted at Nips '24]**: addressed major weaknesses in previous works which generated SAT problems which were either trivial to solve, or generated at such a cost that producing a large dataset for augmentation was intractable. Developed a training procedure for learning models to approximate expensive iterative SAT-algorithms for core refinement.

### Applied Science Intern - Amazon Devices

Seattle, WA

#### MANAGER: VIJAY NIRMALANANDA

Summer 2025, 2026

- Developed a novel, deep model architecture to forecast long-horizon discrete choice within realized and counterfactual markets,
- Extensive statistical and econometric testing, validated by industry experts, demonstrated the model's ability to generalize to complex counterfactual markets and scenarios, including new products, new features, and market-scale price variation such as tariff introduction.

### Associate Researcher (Intern) - Noah's Ark Lab

Montreal, CAN

#### MANAGER: YINGXUE ZHANG

2023 - 2025

- "GraSS: Combining Graph Neural Networks with Expert Knowledge for SAT Solver Selection": joined the project after being blocked for many months. I identified key attributes (permutation sensitivity) in the data which were problematic the model architecture being used at the time, unblocking the project.
- "InnerThoughts: Disentangling Representations and Predictions in LLMs": produced secondary experimental design investigating the role played by individual layers in Large Language Models to answer the question: "Do individual layers encode different semantic/linguistic elements of the text?" We found that certain layers proved instrumental in the LLM's answering of reasoning-based questions.

Professional Experience \_\_\_\_\_

- 2026 **Applied Science Intern**, Amazon
- 2025 **Applied Science Intern**, Amazon
- 2023 **Associate Research Intern**, Noah’s Ark Lab

Publications \_\_\_\_\_

D. Chetelat, H. Zhou, R. Moravej, **J. Cotnareanu**, et. al. 2026. “Statistical Testing of Large Language Models”. **[Submitted to NeurIPS ’26]**

C. Roverato, **J. Cotnareanu**, et. al. 2026. “DynaPFN: Zero-Shot Dynamical System Forecasting with Tabular Prior-Fitted Networks”. **[Submitted to NeurIPS ’26]**

**J. Cotnareanu**, C. Roverato et. al. 2026. “Abductive Reasoning with Probabilistic Commonsense”. Proc. Int. Conf. Machine Learning.

**J. Cotnareanu**, D. Chetelat et. al. 2026. “A Balanced Neurosymbolic Approach to Abductive Commonsense Reasoning”. Proc. Int. Conf. Learning Representations.

F. Regol, **J. Cotnareanu**, et. al. 2026. “Is the acquisition worth the cost? Surrogate losses for Consistent Two-stage Classifiers”. Advances in Neural Information Processing Systems (Spotlight).

**J. Cotnareanu**, Z. Zhang, et. al. 2024. “HardCore Generation: Generating Hard UNSAT problems for Data Augmentation”. Advances in Neural Information Processing Systems.

D. Chetelat, **J. Cotnareanu**, et. al. 2025. “InnerThoughts: Disentangling Representations and Predictions in Large Language Models”. Proc. Int. Conf. Artificial Intelligence and Statistics.

Z. Zhang, D. Chetelat, **J. Cotnareanu**, et. al. 2024. “Grass: Combining graph neural networks with expert knowledge for sata solver selection”. Proc. Conf. Knowledge Discovery and Data Mining.

M. Ozmen, **J. Cotnareanu**. 2023. “Substituting data annotation with balanced updates and collective loss in multi-label text.” Conf. Lifelong Learning Agents.

Awards, Fellowships, & Grants \_\_\_\_\_

- 2022-2024 **Graduate Research Excellence Award**, McGill University
- 2024 **Max Stern Recruitment Fellowship**, McGill University
- 2024 **McConnell Memorial Fellowship**, McGill University
- 2024 **Louis Ho Fellowship**, McGill University

Mentorship Experience \_\_\_\_\_

- |      |                                                                        |                          |
|------|------------------------------------------------------------------------|--------------------------|
| 2025 | <b>Olivier Dionne, Samuel King</b> , Collegiate Summer Research Intern | <i>McGill University</i> |
| 2025 | <b>Deniz Karaku</b> , Undergraduate Honors Thesis Student              | <i>McGill University</i> |
| 2025 | <b>Hongshuo Min</b> , Undergraduate Researcher                         | <i>McGill University</i> |