Calgary ML @ NeurIPS 2024

NEURAL INFORMATION PROCESSING SYSTEMS

UNIVERSITY OF CALGARY

Tuesday, December 10th, 2024

Muslims in ML (MusiML) Workshop

2:30-3:00 p.m. (Lightning Talk) *
4:30-5:00 p.m. (Internal Poster Session), 6:30-8:00 p.m (Joint Poster Session for Affinity Groups)

* A Closer Look at Sparse Training in Deep Reinforcement Learning.
Muhammad Athar Ganaie**, Vincent Michalski, Samira Ebrahimi Kahou, **Yani Ioannou.
This paper explores sparse training in DRL, highlighting methods to improve dynamic sparse training performance at high sparsity, underscoring the need for DRL-specific strategies.

Long-Tail Learning with Language Model Guided Curricula.
Mohammed Adnan**, Rahul Krishnan, **Yani Ioannou.
Improving performance on long-tail classes by leveraging LLMs to build curricula.

Women in Machine Learning (WiML) Workshop

6:30 p.m - 8:00 p.m (Joint Poster Session for Affinity Groups)

Learning to Reweight Examples in Backdoor Defense.
Yufan Feng**, Benjamin Tan, **Yani Ioannou.
We extend the online sample reweighting method from robust learning to the context of backdoor defense.

What's Left After Distillation? How Knowledge Transfer Impacts Fairness and Bias.
Aida Mohammadshahi, Yani Ioannou.
We explore the impact of knowledge distillation temperature on fairness for language and image classification models.

Friday, December 13th, 2024

Main Conference

4:30-7:30 p.m. (Poster Session)

Navigating Extremes: Dynamic Sparsity in Large Output Spaces.
*Nasibullah Nasibullah, Erik Schultheis, **Mike Lasby, Yani Ioannou**, Rohit Babbar.*
Investigates Dynamic Sparse Training for large output spaces. Leveraging semi-structured sparsity, intermediate layers, and auxiliary loss, it enables end-to-end training with millions of labels.
Poster Location: East Exhibit Hall A-C #2004

Saturday, December 14th, 2024

UniReps: Unifying Representations in Neural Models

4:30-7:30 p.m. (Poster Session)

Winning Tickets from Random Initialization: Aligning Masks for Sparse Training.
Rohan Jain, Mohammed Adnan,** Ekansh Sharma, **Yani Ioannou.
Lottery Tickets can't be trained from random init. We show that permuting the mask to align with the new initialization's optimization basin results in a mask that better approaches LTH generalization.