

Colloquium: Machine Learning

27. – 28. February 2025

Showroom (DHEG136E) | Sandgasse 36, Erdgeschoß

It is a pleasure to invite you to the colloquium for our Professorship in Security & Privacy at Graz University of Technology. The public part will be a short educational presentation at Bachelor's level 3rd year in Computer Science on topic Knowledge Graph Embeddings, a scientific talk (titles below), and a discussion with the audience.

Sansone Emanuele

27. February 2025 | 09:00 | Showroom (DHEG136E) | Sandgasse 36, Erdgeschoß

Title: “From Neuro-Symbolic Learning to Unsupervised Deep Induction”

Abstract: Neuro-symbolic learning is an emerging paradigm in machine learning, and more broadly in artificial intelligence, that seeks to integrate data-driven approaches based on deep learning with knowledge-driven strategies rooted in logic systems. The promise of such integration lies in enabling machines to learn and reason from both data and knowledge in a trustworthy manner.

In the first part of the talk, I will discuss the limitations of neuro-symbolic learning, focusing the analysis on two main axes: the amount of human supervision required and the scalability of learning and inference. I will briefly describe some recent advances aimed at addressing these two challenges. Specifically, I will demonstrate how unsupervised representation learning and generative strategies can reduce the need for human supervision, enabling the learning of robust concepts, improving generalization, and lowering data requirements. I will then show how advanced sampling strategies can be used to accelerate learning and inference in neuro-symbolic systems.

In the second part of the talk, I will introduce a grand challenge to solve in the next 5 to 10 years. I will outline a path that goes beyond the current neuro-symbolic paradigm and aims to develop a new generation of machine learning systems, which I refer to as unsupervised deep induction systems. These systems will be capable of acquiring, reasoning, and retaining knowledge from noisy and ambiguous sensory data, similarly to the way humans learn.

Pervez Adeel

27. February 2025 | 11:30 | Showroom (DHEG136E) | Sandgasse 36, Erdgeschoß

Title: “Differential Equations as Neural Network Representations”

Abstract: In my talk I look at the synthesis of neural networks and differential equations. Such models have important applications in the modeling of spatio-temporal dynamics and in inverse problems in science. Past approaches have considered network-in-solver approaches which plug neural networks into classical solvers. I present our work on a solver-in-network approach which allows neural network representations that are differential equations. Embedding specialized, parallel and differentiable solvers in neural networks allows greater flexibility for dynamical modeling. As an application of these models, I discuss our work on the discovery of governing differential equations (ODEs and PDEs) from data.

Özdenizci Ozan

27. February 2025 | 14:30 | Showroom (DHEG136E) | Sandgasse 36, Erdgeschoß

Title: “Resource-Efficient Computing in Adversarial Machine Learning”

Abstract: The rapid growth of digital information is making efficient utilization of data and learning-based computing a shared problem in many domains. Notably, the recent success of machine/deep learning methods in this realm has been accompanied by various technical challenges, particularly regarding their environmental sustainability, and reliability against adversarial behavior. As machine learning systems become an increasingly pervasive technology, it becomes essential to integrate resource-efficiency, robustness, security and privacy into the design of the underlying models, and understand how these aspects are interconnected in challenging ways. In this talk, I will present an innovative research perspective with the goal of answering the question: "How can we sustainably build intelligent machines that can safely and reliably learn to generalize in the open world?". I will give an overview of the state-of-the-art and open problems in the area of adversarial deep learning, and elaborate when and why tailored resource-efficient solutions are necessary in diverse applications ranging from visual computing to natural language processing. I will present our recently developed algorithmic solutions to some of the unique problems in this domain, with an emphasis on the importance of unconventionally exploring the design space of deep learning algorithms from novel perspectives.

Zhu Jia-Jie

27. February 2025 | 17:00 | Showroom (DHEG136E) | Sandgasse 36, Erdgeschoß

Title: “Flow and Transport: Modern Mathematical Foundation for Modern Machine Learning”

Abstract: The theoretical foundation of statistical inference, optimization, and generative modeling has received a major push from the theory of PDE and gradient flows over the last couple of decades. Motivated by several applications in Bayesian inference, optimization, generative modeling, and distributionally robust optimization, I will provide a few state-of-the-art results regarding the Wasserstein and Fisher-Rao gradient flows. These findings not only advance our theoretical understanding but also provide practical tools for enhancing machine learning algorithms. Finally, I will showcase inference and sampling algorithms using a new kernel approximation of the Wasserstein-Fisher-Rao (a.k.a. Hellinger-Kantorovich) gradient flows, which have better convergence characterization and improved performance in computation.

Kappel David

28. February 2025 | 08:30 | Showroom (DHEG136E) | Sandgasse 36, Erdgeschoß

Title: “Biological principles for efficient machine learning”

Abstract: Recent advances in machine learning (ML) have demonstrated impressive performance on complex tasks such as human-level image understanding and natural language processing. These ML models rely on artificial neural networks that, like biological brains, use billions of neurons and synapses to process complex stimuli. However, unlike biological brains, these neural networks consume vast amounts of energy, with a single training session often exceeding the energy and carbon footprint of a car over its lifetime. If the current rate of growth continues, ML models could overtake the transport sector in the global energy balance within 10-20 years, posing another major threat to mitigating climate collapse. In this talk, I will highlight the mechanisms that enable the amazing energy efficiency of biological brains. Based on these findings, I will present new approaches

to significantly reduce the energy footprint using hybrid ML/bio-inspired models.

Bellec Guillaume

28. February 2025 | 11:30 | Showroom (DHEG136E) | Sandgasse 36, Erdgeschoß

Title: From neuroscience to hardware-first machine learning algorithms

Abstract: The increasing size of AI models is calling for alternative computing principles. We advocate that computational principles from neuroscience continue to provide inspiration for new computing principles. Looking back at the data on synaptic plasticity from the last decades, we present theory-driven learning algorithms which are designed for contemporary energy-efficient hardware. These neuromorphic algorithms alleviate multiple features of back-propagation learning which are incompatible with a physical always-on learning hardware. Looking forward, we then speculate how the next generation of brain-inspired computing principles can be derived from contemporary brain recordings. Using a new network reconstructions technique we show that it will become possible to measure "brain gradients" and refine our algorithmic understanding of learning in the brain.

Stoian Mihaela

28. February 2025 | 14:30 | Showroom (DHEG136E) | Sandgasse 36, Erdgeschoß

Title: "Refining Deep Generative Modelling using Background Knowledge"

Abstract: Synthesising realistic tabular data often relies on deep generative models. However, these models fail to account for inherent relationships between features, encoded as background knowledge, which synthetic samples must satisfy to be deemed realistic. Existing methods handle non-compliant samples by discarding them, leading to potentially indefinite inference times. In this talk, I will present a novel approach that integrates a constraint layer directly into the topology of deep generative models to account for the relationships between the features. The layer automatically incorporates background knowledge requirements and ensures compliance with these constraints during both training and inference. I will first present a method for handling linear constraints and then discuss its recent extension to support constraints as expressive as disjunctions over linear inequalities, capable of modelling non-convex and disconnected spaces. The layer not only guarantees that such constraints are satisfied, but also significantly improves the machine learning efficacy of deep generative models without hindering sample generation times. I will conclude by discussing how this framework contributes to the broader vision of bringing neuro-symbolic AI onto the stage of real-world applications.

Li Jun

28. February 2025 | 17:00 | Showroom (DHEG136E) | Sandgasse 36, Erdgeschoß

Title: "Test Time Training with Generative Models"

Abstract: Test Time Training (TTT) has emerged as a promising research direction in machine learning, offering a paradigm shift in how models are deployed. Unlike traditional approaches where models are trained offline and then used as fixed entities, TTT allows for further fine-tuning of a model's parameters specifically for each test instance or a small batch of instances without labels. This dynamic adaptation can enhance performance and adaptability, particularly when combined with the generative capabilities of modern models. Our research investigates the synergy between TTT and generative models, exploring how this combination can unlock new levels of performance

and flexibility in various applications.